

SPECIFICATION

Docket No. 0635MH-40874

TO ALL WHOM IT MAY CONCERN:

BE IT KNOWN that we, Frederick S. M. Herz, David C. Parks, and Sampath Kannan, residing in the state of Pennsylvania, Paul Labys, residing in the state of Utah, and Jason Eisner, residing in the state of New York, have invented new and useful improvements in a

SECURE DATA INTERCHANGE

of which the following is a specification:

CROSS REFERENCE TO RELATED APPLICATIONS

1 The present application claims the benefit of United States Provisional
2 Application No. 60/161,640, filed October 29, 1999, titled SECURE DATA
3 INTERCHANGE, and Provisional Application No. 60/206,538, filed May 23, 1999,
4 titled SECURE DATA INTERCHANGE, both of which are hereby incorporated by
5 reference.

BACKGROUND OF THE INVENTION

6 1. Field of the Invention:

7 The Secure Data Interchange invention describes a system to allow a privacy-
8 protected market for data exchange between multiple self-interested parties. The system
9 presents a general infrastructure for the exchange of information within a safe privacy-
10 protected environment, between multiple self-interested parties. We propose a central data
11 warehouse that maintains data submitted by different users, and executes queries and
12 programs on the data. Rules are associated with data that define how the data can be used
13 and queried, to allow agents that submit data to maintain absolute control over its use. SDI
14 acts as a trusted-intermediary to all parties, and implements an internal market for queries on
15 the information, allowing agents to specify prices for data access. Furthermore, SDI
16 supports complex queries such as collaborative filtering, that can provide a querying agent
17 with a one-time benefit of data access but without long-term access to the data that was used
18 to compute valuable results.

19 The invention relates to systems that provide personalized information, profiling,
20 automated matchmaking and information exchange, providing a framework that protects
21 privacy and allows information collection and profiling within a carefully controlled
22 environment. Although the marginal cost of data duplication is small, there are hidden costs
23 associated with data, for example because of privacy concerns, and data can be a valuable
24 resource in many problems. In business-to-consumer (B2C) applications Secure Data
25 Interchange addresses the direct conflict between the goal of personalization and the need
26 for privacy, preventing the exchange and collection of information without knowledge and

1 consent. In business-to-business e-commerce applications (B2B) SDI allows vendors to
2 provide sensitive and valuable information, for example about business needs and customer
3 bases, in a secure environment that controls access and leverages value.
4

5 2. Description of the Prior Art:

6 The invention of Secure Data Interchange relates to a wide-range of application
7 domains, all of which are characterized by a need to exchange information in a privacy-
8 protected and carefully controlled market-based environment.

9 As a key application we suggest a system for personalized information delivery in a
10 networked environment, in which the SDI-proxy can fact as a local filter on information,
11 based on what it knows about a user's preferences and methods for filtering pushed by the
12 provider of content. The system allows collaborative filtering through information that is
13 provided to the central data warehouse, but never released directly to other agents;
14 collaborative filtering methods are computed in the central SDI data warehouse. Further
15 motivation is provided with reference to some electronic commerce applications, that we
16 describe in (A) business-to-consumer and (B) business-to-business e-commerce
17 applications.

18 In addition to applications within commerce, the system of Secure Data Interchanges
19 is central to developing many other new products. Examples include the formation of "self-
20 help" groups between a set of individuals with common interests, and applications to
21 personal information delivery systems, e.g. for educational and informational purposes.
22

23 A. Business-to-Consumer (B2C) Electronic Commerce.
24

25 The recent explosion of electronic commerce, in particular Internet-based individual-to-
26 business electronic commerce, presents new opportunities for automated personalized
27 information delivery and the automated customization of products and services. This type of
28 personalization is very valuable to vendors because it can increase sales volumes, enable
29 cross-selling and up-selling of goods and services, and allows vendors to price products

1 dynamically based on information about the preferences and goals of customers.
2 Personalization is also useful to customers when it correctly identifies the requirements and
3 preferences of a customer, because it can reduce search cost and enhance the “shopping
4 experience”. Perhaps a customer can find the good or service (i.e. desirable
5 price/quality/feature tradeoff) that he/she wants more quickly than without personalization,
6 or receive information about an interesting new product or service that he/she did not know
7 about.

8 The basis for these new services is that Internet-based “shop fronts” can be
9 individualized on a per-customer basis, dynamically and in real-time. Traditional main-street
10 shops must offer the same store layout to every customer, because the layout is physical,
11 although some level of personalized service can be achieved through well-trained sales
12 assistants, that act as a “guide” for a customer within a store. On-line “shop fronts” are
13 virtual, and configurable at negligible cost to the customer or the vendor, assuming that
14 computation is cheap and fast.

15 Furthermore, Internet-based electronic commerce can allow business to collect vast
16 amounts of consumer information, because customers interact through a computer-based
17 interface. Customers can be monitored as they browse a Web site for products and services.
18 Information such as the search-terms that users enter into a search engine, the links that
19 users follow, and the length of time spent on each page, can all provide an insight into the
20 current goal of a customer, i.e. the type of product that he/she wants. When combined across
21 different sessions, and with similar information about the browsing and purchasing habits of
22 other customers, the information can be folded into a long-term view of the preferences and
23 needs of a customer.

24 Moreover, new network connectivity enables different vendors to exchange profiles
25 for common customers, either statically or dynamically, in order to build broad and detailed
26 profiles across vendor domains. There exist many potentially powerful synergies between
27 the data sets that are collected by different vendors, that can be leveraged to provide
28 appropriate services and products to customers. When analyzed with the proper statistical
29 tools these data sets can reveal fundamental patterns in the behavior of users, and enable a
30 vendor to provide appropriate information to a user. Furthermore, access to user-profiles

collected by other vendors can enable vendors to provide focused information delivery to first-time users, and also cross-market services with other appropriate vendors.

Providing user profile information within a carefully controlled environment can benefit vendors and users:

- Vendors would find benefit in sharing data with other vendors; this would deepen their understanding of their customers' behaviors and preferences, especially if some customers were traceable across several data sets.
- Users would benefit from sharing data with other users. This is already evident in the popularity of news groups and web discussion pages catering to individuals with shared interests. By learning what other people with similar tastes and preferences have discovered and enjoyed, a user can sidestep information overload in the search for personally satisfying information.
- Vendors can benefit from receiving data about users. An obvious example would be in the use of collaborative filtering for the marketing of targeted promotions; rather than being deluged with coupons and advertisements that are of absolutely no interest, a user would benefit by being presented with advertising that is highly relevant. In the process, the vendor would increase advertising response rates, boosting overall efficiency.
- Users can receive benefits from providing information to vendors. Personalization of content at vendors' web pages, and well-focused banner advertisements at other web sites that they visit.

The problem is that a user wants controlled personalization, in the sense that it might not be desirable for information about every on-line transaction that a user performs, every on-line document that a user reads, and every web page that a user visits, and demographic information, to be available to every business that the user interacts with, in the virtual and physical world.

A.1 Focused Banner Advertising/Content provision

Internet-based media sites have followed preceding formats in generating revenue from advertising, with content to users often provided free-of-charge. The business model is

1 similar to that in newspapers, magazines, and television, where circulation and
2 audience/readership demographics are used to drive revenue. Electronic media presents new
3 opportunities for media-based business: for example multimedia techniques and
4 interactivity, personalized delivery of information, and personalized targeting of
5 advertising.

6 The problem – as before, is to acquire and leverage information about the
7 preferences and interests of a user, within a system that protects user privacy (i.e. controls
8 the collection and exchange of information about users, and controls the use that is made of
9 that information). A further problem is to extrapolate information from a large corpus of
10 data about an individual user.

11 12 **A.2 Mailing Lists**

13
14 As another example, suppose that business A requests a list of individuals that meet a
15 particular criteria. Consumer B meets the criteria, but is only listed for business A if A also
16 meets criteria specified by B, for example if A will provide information about new products
17 and services that are interesting to B. In an application to the profiling of users on-line, the
18 problem is that users want to receive the benefits of targeted products and advertisements,
19 but want to avoid the abuse of profile information and control vendors' access to that
20 information.

21 22 **B. Business-to-Business (B2B) Electronic Commerce**

23
24 The Internet provides businesses with network connectivity with other business, both
25 competitors and partners. This connectivity allows businesses to exchange information
26 about customers (dynamically or statically), in order to identify potential new customers,
27 build better profiles for existing customers, and up-sell/cross-sell products and services in
28 real-time. The problem with this exchange of information (that can include swaps, sells, and
29 rental access) is that businesses need to (a) protect the privacy of their customers; (b)
30 prevent information release to competitors, either directly or through third-parties.

B.1 Privacy-Protected Identification of Synergies/Matches

There are many scenarios where autonomous agents would like to be informed of matches under conditions of mutual consent, but without information leakage to any agent if any one of the agents declines the match. Consider two vendors, A and B, and suppose the vendors seek strategic partnerships with other vendors that have appropriate skills and goals. However, vendor A does not want to broadcast to all vendors its need for a business partner or a new alliance, instead vendor A wants to be introduced to another vendor with the right mix of capabilities; similarly for vendor B. What is required is a system that only introduces vendor A to vendor B, and perhaps anonymously at first, if both vendors consent to the introduction. The problem is to provide information that enables matches, without allowing bad matches and abuse of information – i.e. within an environment of secure data interchange.

B.2 Credential-based Introductions, Contracting and Messaging-systems.

There are many situations where individual parties, for example individuals or businesses, require introductions to credentialed individuals and/or businesses, with the aim of building a new relationship or making a new contract. Consider for example business associations, where credentials about non-bankruptcy, and no previous attempts to defraud could be important. Consider social introductions, where individuals might be concerned about past criminal activities of new contacts. In the domain of automobiles, we could consider a system that identifies other automobiles in the physical location of a vehicle that have recently been involved in an accident. The problem is to manage certificates within a system where users can maintain multiple identities, and to protect the release of certificates without suitable provisions for terms-of-use and criteria for request.

SUMMARY OF THE INVENTION

1 The above problems are solved, and a technical advance achieved, by the system of
2 Secure Data Interchange. The Secure Data Interchange system enables information about
3 bilateral and multilateral interactions between multiple persistent parties to be exchanged
4 and leveraged within an environment that uses a combination of techniques to control access
5 to information, release of information, and matching of information back to parties.

6 The system of Secure Data Interchange (SDI) provides a trusted server containing
7 a large database of information that is owned by its providers. Each data record has an
8 associated price rule, that controls access to data. The pricing model allows a data owner
9 to specify a price for different types and amounts of information access, and whether the
10 identity of the information owner is required, and the system of SDI computes a
11 composite price for a query based on aggregated prices for a query over a number of
12 different data owners, with an internal market that favors low priced data. The pricing
13 model allows discounts based on certificates of a requesting agent, and as a special case
14 implements the standard capability-based access control systems, where information is
15 provided to users with appropriate permissions (i.e. with zero and infinite prices). In
16 addition, the system of Secure Data Interchange allows data to be submitted with a level
17 of random perturbation (noise), to provide added privacy protection, or alternatively
18 allow an agent to specify in conditions under which additional noise should be added to
19 data. A query is priced before execution, to allow an agent to decide whether or not to
20 execute a query, and select between alternative types of queries. Binding price quotes are
21 provided to querying agents, and queries can be scaled to meet a budget.

22 Data owners can submit data to the central data warehouse with different degrees
23 of identification, for example anonymously, pseudonymously, or with a true identity. For
24 example, in the case of data that represents a user's profile information, for example
25 information about the interests of a user, a user might prefer to use a number of different
26 pseudonyms for different types of activities that he/she likes to engage in online. A user
27 might maintain a number of different aliases within the database, for example to represent
28 different types of things he/she likes to do which have little bearing on each other.

1 manages a user's interactions with the central SDI data warehouse, i.e. providing profile
2 information and controlling profile access. The client-side SDI proxy for an agent that
3 represents an individual browsing the Internet can manage that user's profiles in
4 interactions with other agents, for example representing vendors and content providers.
5 The client-side SDI proxy can also handle decisions about what types of information to
6 submit to the server, and manages query execution on behalf of the agent. The client-side
7 SDI proxy agent can also push information about a user's on-line activities to the central
8 SDI data-warehouse in real time. This enables a system of "time-of-purchase-
9 competition" system, in which a user can request competitive counteroffers from other
10 vendors before making a purchase.

11 The system addresses the fundamental conflict that exists between rights of privacy
12 and efficiency gains from better bilateral exchange of profile/preference information. SDI as
13 applied to B2C e-commerce allows consumers to receive targeted information about
14 products and services, but without the loss-of-privacy that can easily occur in the current on-
15 line profiling "free-for-all". The cookie technology provided by Netscape to supported
16 personalized sessions with a single vendor on-line has been used by advertising network
17 providers such as DoubleClick to track users across multiple sites, often without either the
18 consent or knowledge of that individual [New York Times, Feb 7, 2000].

19 In describing the system of secure data interchange we claim the following novel
20 technical ideas:

- 21 (a) Agents can associate price-rules with information that is placed in the central
22 database, and retain absolute control and ownership over all uses of that data.
23 As a special case of price rules, the system supports access based on certified
24 properties of querying agents (with zero and infinite prices). The pricing
25 model allows information providing agents to receive direct value for data,
26 and allows agents that request access to information to receive a price before
27 a query is executed, and make appropriate decisions about what type of
28 queries to execute.
- 29 (b) A number of novel techniques are proposed to allow data processing within
30 the data warehouse without releasing too much information to an agent.
31 General programs, for example collaborative filtering techniques, can be

executed in situ within the data warehouse, so that agents can receive the aggregate benefits of information, without receiving details about the information. We might also allow an agent to receive anonymous or randomized information, and control the amount of information that can be received from a particular record.

(c) The system of SDI can act as a trusted intermediary between agents, notifying agents about information of a particular type, and for example sending messages between agents without breaking the identity of agents except by consent. A special type of query that we call a persistent query allows an agent to maintain a permanent “searching” presence in the central database, always on the lookout for useful information.

(d) In a key variation we push a small version of the SDI data warehouse onto an user’s client computer, where it acts as a proxy agent, and further proxies an agent’s interaction with other agents in real time. The local SDI data warehouse, called the client-side SDI database, is a trusted party to the client, trusted to maintain information that it is valuable to the client and use that information in appropriate ways. The advantage of this method is that an individual never needs to release sensitive profile information, it is always held on its local computer—but can still receive the benefits of personalization.

(e) We also suggest a client-side SDI proxy that can collect information about a user, for example within an Internet browsing application, and periodically push the collected information to the SDI data warehouse in a controlled way. The client-side SDI proxy can also be responsible for certain data certification functions, and can manage a user’s interactions with other agents to protect its privacy in non-SDI mediated transactions.

As an application to B2C e-commerce, the system of SDI allows client-side personalization instead of provider-side personalization. Instead of passing profile information to a provider and receiving personalized information in return, providers can provide personalization methods that are used interactively with local profile information about consumers to target products and services without receiving explicit information

1 about a user's profile. In a simple form, the vendor provides complete information about
2 its services, and a method to display them to the user based on his/her local profile. When
3 describing the application of SDI to electronic commerce we also describe methods to
4 implement necessary ancillary systems that are essential to supporting full e-commerce
5 functionality within an identity-protected system, such as systems for pseudonymous
6 payments and physical mailing of products.

7 Collaborative filtering based on profiling information from multiple users is
8 supported within the central SDI data warehouse, but within a system of economic
9 incentives, where users provide profile information in return for receiving payments from
10 vendors for that information. This allows broad network-wide information to be used for
11 profiling, in addition to deep vendor-specific information.

09699098-102700

BRIEF DESCRIPTION OF THE DRAWINGS

1 The novel features believed characteristic of the invention are set forth in the
2 appended claims. The invention itself however, as well as a preferred mode of use, further
3 objects and advantages thereof, will best be understood by reference to the following
4 detailed description of an illustrative embodiment when read in conjunction with the
5 accompanying drawings, wherein:

6 Figures 1 through 21 illustrate various parts and embodiments of the invention.

002207-3606960

DESCRIPTION OF THE PREFERRED EMBODIMENT

1. Introduction

The invention of Secure Data Interchange (SDI) describes a general infrastructure for the exchange of information within a controlled environment. We propose a central data warehouse that maintains data submitted by different users, and executes queries and programs on the data. Rules are associated with data that define how the data can be used and queried. As such, the system of SDI prevents the exchange and collection of information without knowledge and consent. The system allows for payments to be received by the providers of information, in return for data access. In application, the invention enables new systems for the delivery of personalized information, profiling and automated matchmaking and information exchange, all within a framework that protects privacy and maintains data security. SDI supports the collection and exchange of information between, and relating to, autonomous (and possibly self-interested) agents within a distributed environment.

1.1 Definitions.

Agent. An agent in SDI is a any party that wants to sell or give away data to other parties, or buy or receive data, or in general both provide and receive data. Agents may represent any party with individual goals, autonomy of control, and a persistent identity. Examples in business-to-consumer e-commerce include business such as newspapers, book stores and travel companies, that wish to receive data about the profiles and buying habits of users so that they can personalize the information, products and services that are sold to users. We assume that agents are autonomous from the system of Secure Data Interchange, and follow actions consistent with their preferences, abilities and resources. Similarly, an on-line consumer is an agent that wishes to provide data about its preferences and buying habits to vendors and other consumers, in return for well targeted products and financial reward.

Agent Computers. Agents are represented in the system of secure data interchange with dedicated computational resources, agent computers, with permanent memory, processing power, and network connectivity. For example, an agent computer might be a consumer's home PC that will act as a client machine in interactions with the central SDI data server and

1 the servers of on-line vendors. For a vendor, an agent computer might be the server
2 computer that it uses to execute its on-line business. Agent computers might also be thin
3 clients, such as mobile computing devices, handheld devices, cell phones. We push different
4 amounts of data and functionality within SDI to agent computers and away from web centric
5 devices depending on the nature of the computer and an agent's preferences.

6 **Certificates.** Certificates are used within SDI to establish trust between different agents,
7 and to help agents to reach useful agreements.

8 **Profile:** Each agent can have one or more profiles, which the agent can assume in its
9 interactions with other agents. Part of a profile is a user's **identity**, that provides a (possibly
10 limited) method for another agent to identify the agent in the future when it assumes the
11 same profile. We allow three types of agent identities: anonymous, pseudonymous,
12 persistent pseudonymous, and true identity. An agent may assume a profile and an identity
13 in any interaction with another agent.

- 14 • **Anonymous.** To assume an anonymous identity agent A creates a one-time identifier
15 that it uses in interactions with one other agent. The identifier may allow the other
16 agent to respond zero or one times.
- 17 • **Pseudonymous.** A pseudonymous identifier created by agent A may be used with
18 more than one other agent, and allows agents to respond to agent A as many times
19 as they like, but agent A can terminate the pseudonym at any time and separate
20 from the identity. Agent A can optionally restrict the number of agents that may
21 reply to the pseudonym, for example to the agents that it explicitly provides with the
22 identifier.
- 23 • **Persistent pseudonymous.** A persistent pseudonymous identifier is created by agent
24 A for use with agents in set S, and provides the added condition that agent A
25 promises to use the same pseudonym for all interactions with all agents in set S for
26 all time into the future. Agent A can optionally restrict the number of agents that
27 may reply to the pseudonym, for example to the agents that it explicitly provides
28 with the identifier.
- 29 • **True identity.** If agent A interacts with another agent under its true identity then it
30 has no method to prevent the agent or any other agent responding to agent A in the
31 future.

1 The identity that an agent assumes with a profile matters not only to the ability of an
2 agent to control the agents that can send it messages, but also its ability to control the
3 amount of information that can be exchanged about the agent in the open marketplace, out
4 of the agent's control. A basic premise in the system of Secure Data Interchange is to keep
5 control of data. We provide methods that allow an agent to release data, or performing
6 processing on data, to one agent but prevent that agent from selling the data on to another
7 agent with which agent A also interacts. For example, one key technique is for agent A to
8 use a unique pseudonym with every agent that it interacts. This allows an agent (so long as it
9 is careful not to release other identifying information) to release profile information to other
10 agents without losing the value of that information, because there can be no secondary
11 market in the agent's profile. If agent A provides information to agent B, then agent B
12 cannot pass that information onto another agent C and have it still linked to agent A because
13 agent C does not know the identity of agent A among the agents that it interacts with.

14 An agent's profile contains any and all data that an agent might wish to exchange
15 with another agent when it assumes a particular identity. However, just because the data is
16 in the profile it does not mean that it is available to another agent. Possible information in
17 the profile of an on-line consumer includes: transactions that it has performed with other
18 agents; information that relates to its true identity (e.g. salary range or education level);
19 information provided by the agent (such as its preferences for a particular type of product,
20 etc.); and other information that has been compiled based on observing the behavior of the
21 agent (e.g. physical location for a mobile user, such as a user in a vehicle, or trace of recent
22 web pages visited for a user that is browsing the Internet.)

23 **SDI-Proxy:** An SDI-proxy refers to the software that runs on top of an agent's computer
24 device, and configures that device for Secure Data Interchange. The proxy intermediates
25 interactions between pairs of agents, and also intermediates transactions between agents and
26 the SDI data server. For example, an on-line consumer might define a profile-management
27 policy at the client-side SDI proxy that automatically configures the agent's profile and
28 identity when as the agent interacts with other agents. The profile-management policy
29 implements an appropriate policy to select the user's profile and identity on the basis of the
30 information that is available about the other agent (for example from certificates).

1.2 System Architecture

The basic architecture for SDI is a system of agent computers, connected via a network (Internet, wireless, or otherwise) to other agent computers, and with a central SDI shared data warehouse. In Figure 1 we illustrate the top-level architecture of Secure Data Interchange, the networked system of agent computers and a centralized server computer that acts as a repository for data, rules and code. This is called the SDI data warehouse. Each user is associated with an agent computer, and in general users can be individuals, groups of individuals, or companies. In its most general form, the system of Secure Data Interchange is for a system of multiple autonomous agent computers, involved in multilateral communication. We restrict the system description to bilateral communication between agents, without loss of generality because any multilateral (multicast or broadcast) can be implemented as a set of bilateral communications. In any bilateral communication there are two parties, the sender and the receiver (and the parties can dynamically change over a communication session).

SDI proxy agents and profile management policies on agent computers mediate bilateral agent interactions. Agents define profile-management policies that are implemented on agent computers and determine appropriate profiles and identities for an agent in interactions. Every time an agent initiates a new interaction, new information is available about the agent, that can be compiled by the agent's own computer device, and also by the computer device of the agent with which it interacts. Careful profile and identity management provides an agent with absolute control over the ability of other agents to profile the agent and exchange information about the agent, for example the agent can use a unique pseudonymous identifier with every other agent. However, the core of the SDI invention is that we encourage agents to exchange information, by providing a secure central data interchange for that purpose. Agents can submit profile information, and other data, to the central data warehouse and make it available for particular types of data mining by other agents, and receive financial reward for providing data. Products and services can also be targeted for agent A on that agent's computer, without the provider agent receiving information about agent A's profile. For example, vendor B might provide generic information about its products to agent A, and a personalization rule that it has constructed

1 from data mining the central SDI data warehouse. The agent computer of agent A can then
2 use its local and private profile information to decide what products to target to the agent,
3 based on the instructions that it receives from agent B.

4 The arrows indicate possible flows of information between the different elements
5 within the system. We allow agents to communicate directly, via bilateral or multilateral
6 communication, and also indirectly via the central data warehouse. Information can be
7 submitted to the central data warehouse, for example new data and new rules, or queries
8 that the data warehouse will execute. The data warehouse returns information to agents, for
9 example the results of a query. The Secure Data Interchange invention is independent of the
10 implementation details of the communication platform. We assume that the system of SDI is
11 built on top of a secure communications platform, for example via a SSL-encrypted TCP/IP
12 session in an application to the Internet. Furthermore, we draw on cryptographic techniques
13 known in the art for identity management, and additional techniques to support
14 pseudonymous identities within a particular network protocol, e.g. the stripping of sender
15 URL addresses from HTTP packets on the Internet.

16 In overview, the top-level claim in the system of Secure Data Interchange, of a
17 privacy-protected market-based system for data exchange between self-interested parties, is
18 constructed from the following core components:

- 19 ○ A Central Data Warehouse, with associated price-rules and constraints submitted by
20 the owners of data. Architectural variations allow virtual links to data in the central
21 data warehouse, with data physically located in distributed locations.
- 22 ○ A Query-execution and Price module, which executes queries on the data within
23 constraints on the types of queries that can be performed, and computes the price of
24 queries, collecting payment.

25 and the following optional components:

- 26 ○ Distributed data management, via client-side SDI “proxy” agents, that intermediate
27 the interactions between agents, and manage data provided to the central SDI data
28 warehouse. Distributed data management includes methods to manage an agent’s
29 identity in interactions with other agents, for example via pseudonymous and
30 anonymous interactions.

- 1 ○ Distributed query-execution, via client-side SDI query-execution modules, that
- 2 allow data that is physically stored in distributed locations to be queried on
- 3 distributed nodes, without an explicit release of the data.
- 4 ○ Data perturbation methods to augment data values with randomized noise, to allow
- 5 queries to tradeoff price and quality, and to allow a user to protect his/her identity by
- 6 hiding revealing details of submitted data elements.

7 In describing key applications of the system of Secure Data Interchange, we describe
8 in later sections additional features that extend the functionality of the invention:

- 9 ○ Community dollars: an extended payment scheme that allows extended forms of
- 10 payment within the system, for example allowing payment to be made in terms of
- 11 discounts in purchases from particular vendors, or in return for agreements to
- 12 purchase a certain number of products over an extended period of time.
- 13 ○ Within an Internet-browser based system, methods to control a user's browsing
- 14 experience via a personalized portal, where personalization is performed via data
- 15 mining techniques executed by the system of SDI on data in the central data
- 16 warehouse.
- 17 ○ Methods to support anonymous and pseudonymous electronic commerce, e.g.
- 18 delivery and payment services.

19 In the next section we provide an overview of each of the central SDI data warehouse,
20 which forms the core component of Secure Data Interchange. We then add technical
21 details for specific components, to make our ideas more concrete.

22 1.3 Implementation Details

23 In this we describe some of the core technologies known in the art that would be used to
24 build up an implementation of the SDI system. The technologies span areas in
25 cryptography, for pseudonym management, digital certificates, payment mechanisms,
26 etc.; and information theoretic-methods, for example to protect the identity of the
27 originator of a message by routing messages through local "crowds" of agents.

28 1.3.1 Certificate Management

Cryptographic techniques well-know in the art [Chaum 81; Chaum 85; Chaum 91] provide the ability for a certificate to be linked to an agent's identity, and not transferred to other agents. This is the basic functionality required of a certificate management system. For example, a certifying agency can sign the public key of an agent with the private key of the certifying agency, to indicate that the agent satisfies requirements for certification. Another agent can verify the certificate with the public key of the agent and the public key of the certifying agency. The certificate cannot be transferred to another agent unless that agent assumes the same public key. We assume a public key infrastructure to manage this process.

1.3.2 Support for Anonymous and Pseudonymous Identities

We have already noted that it is important to provide whatever additional support is required because of the underlying communication infrastructure to protect agents' profile management policies. For example, in the TCP/IP mechanism a message must be stripped of the network address of the originating Internet server, because this can provide information to allow pseudonyms to be linked. Similarly, messages can be routed through a common gateway or random "forwarders" as in the "CROWDS" system [GGMM98; RR98] to provide pseudonymity. Furthermore, other e-commerce functions, such as payment and the anonymous mailing of goods must be supported (see [INSERT A FORWARD REF TO A LATE SECTION OF BOTTOM-LEVEL SDI DOC]).

The ability to embed data within web pages allows client-side processing of information. By embedding profile and location information directly within a web document we can alleviate the bandwidth and computational bottlenecks that can occur at a centralized profile server if profiles are fetched on-the-fly when web pages are downloaded by clients. The origin server (supported by the vendor) requests periodic profile updates from the central SDI server. This duplication of information enables the profile and the page contents to be provided directly from a vendor's server.

There are some potential drawbacks of this approach: (1) the profile information associated with a web page and target objects can be out-of-date; (2) the profile information is available to all clients and proxy servers, not just those that are SDI-enabled; (3) the profile information can be altered. We suggest technical solutions to each of these problems below.

1 In addition, we prevent unauthorized access of embedded profile information
2 through the encryption of the metadata that is represented within the XML structure of a
3 web page.

4 Profile information can be encrypted using a hierarchy of keys, so that different
5 levels of access to the information may be provided according to the access levels of users
6 and vendors. All users that request web pages from SDI-enabled vendors, whether or not the
7 user is a member of SDI receive the same profile information. We provide encrypted
8 profiles to vendors in the 'Profile Update' messages from SDI to vendor servers, so that: (a)
9 unauthorized agents cannot tamper with the profiles; (b) the profiles cannot be read by
10 unauthorized agents.

11 The SDI system supplies a private key to trusted SDI client software, that enables
12 only SDI-enabled clients to access profile information, and only access that information to
13 the extent permitted by privacy policies of users and vendors. Different levels of encryption
14 enforce multiple levels of access. Periodically the key pairs are changed to prevent
15 extended attempts at cryptographic attacks. The SDI system uploads the key that provides
16 the correct level of access for a user to a user's client, once terms of access and profile
17 management have been agreed. A client can only access embedded information once
18 enabled with a relevant key. Finally, profile information is signed with a digital certificate,
19 to prevent third parties from tampering with profiles for commercial gain.

20 2. Core Modules

21

22 Figure 2 illustrates the core modules within the system of Secure Data Interchange. In
23 this section we briefly describe each module in relation to the other modules, and provide
24 more details in later sections.

25 The central SDI data warehouse is the core system in SDI, managing data records on
26 behalf of agents. In combination with the query-execution module, these two modules
27 implement the privacy-protected market for query-execution, where queries are executed:

- 28 • If the querying agent has the right certificates, as defined by the owner of data
- 29 • (and) If the querying agent pays the cost of performing a query

1 syntax of valid queries on a particular type of data, and executes queries. As data types
2 become standardized, we could suppose a library of standard accessor functions for
3 different data types, perhaps provided by third parties. For example, a special data type
4 could be developed to represent the profiles of online customers. As part of query
5 execution therefore, is the execution of methods associated with heterogeneous types of
6 data. Of course, we can also describe a simple special case where the data types are more
7 transparent to SDI, and data fields are directly indexed as in standard databases.

8 Agents can submit queries to the SDI data warehouse, these are placed in the pending
9 queries queue if they are one-time and to be executed as quickly as possible, or in the
10 persistent queries queue if they are to be executed periodically, as the data in the
11 repository changes. The query execution module contains the processor that performs
12 instructions on behalf of agents, within the constraints of price rules.

13 Price rules are associated with data records. Price rules compute a charge for
14 requested information about the data record. We describe a number of simple price
15 languages, to allow an owner of information to configure access to the data. In general
16 the price rules can be quite complex, and implemented within an object-oriented
17 framework, so that the system of SDI is expandable to new data types. The query-
18 execution module can execute both the methods to access data in objects, and the
19 methods to price access to data. In performing query execution, we suggest some
20 techniques to minimize the price paid to perform a query—for example we implement a
21 simple internal marketplace where the data records with the smallest price are selected,
22 with all other things equal.

23 As a special case, prices can implement the standard capability-based access control
24 systems, such as those in UNIX, where a user must possess the necessary properties to
25 read information. A price \$0 corresponds to permission to access information, while a
26 price \$infinity prevents a user from reading information. The price rule computes a price
27 for accessing information based on a number of factors: the properties of the requesting
28 agent, as demonstrated with certificates; the information requested; and whether the
29 identity of the owner of the information is revealed to the requesting agent. SDI enforces
30 the price rules, with payment collected from agents before the results of queries are
31 reported. SDI allows an agent to contribute information to the shared database but

1 against each record pointer, and relate to the degree of information provided about that
2 record. For example, consider the problem of finding a set of profiles of a particular type.
3 As the search is performed over data records in the database the query-execution module
4 assesses a provisional payment, as records are accessed. However, the requestor of the
5 information only pays the sum payment charged by the records that are actually selected
6 as suitable. The owners of the data records that were not selected in the final result
7 returned to the user receive no payment, because no information about their data was
8 released and the information was only used in intermediate processing steps. This is
9 important, because it makes it unnecessary to formulate queries carefully in order to
10 avoid extra cost because of redundant intermediate steps.

11 Complex queries, including query-programs, with intermediate results and
12 anonymous record pointers expand the space of data manipulation in a system with
13 providers of information that are more willing to release information if that information is
14 never made directly available to another agent except in aggregated system-wide terms.
15 An agent might be willing to allow submitted data to be used for data mining applications
16 so long as the data remains in the trusted data repository. The central SDI server can
17 provide standard types of query-programs, stored in the SDI-methods module, for
18 example methods to perform collaborative filtering.

19 The query-execution module also performs persistent queries that are submitted by
20 agents to be executed when particular conditions in the data repository are true. Persistent
21 queries check for certain information to reside in the data warehouse, and notify a
22 requesting agent whenever these conditions exist, perhaps automatically making an
23 action.

24 The pricing module is responsible for pricing queries before execution. The basic
25 problem is to aggregate the charge over all data records that provide information to form
26 part of the result of the query, and to estimate the price of a complex query without
27 performing the query. We suggest a simple top-level query language to allow a user to
28 price quality-cost tradeoffs in the query that is finally executed. For example, statistical
29 techniques can be used to compute aggregate statistics without accessing all data records
30 in the database. To give another example, a query can be performed on the basis of
31 accurate information or on the basis of slightly randomized information. The pricing

1 to provide “translation services” to convert local data formats into the SDI native data
2 format. A typical technology to support a shared ontological representation is XML
3 (Extended Meta Language), which allows a grammar to be defined for a document, with
4 meaning embedded in tags. The trend towards XML-based applications should facilitate a
5 shared ontology structure, and allow metainformation to be associated with information
6 and describe data. XML allows intelligent integration of data from multiple databases.

7 Alternatively, we can allow data to be stored in heterogeneous formats across a
8 single unified database, within an object-oriented infrastructure. Each data object has a
9 “wrapper” that controls access, and provides an interface for queries. When a query is
10 executed, the method is invoked, and the result computed with the method and the data.
11 This is perhaps more efficient than the aforementioned approach, in that efficient data
12 formats are query structures are retained.

13 3.1 Data Structure

14 In this section we describe the structure of the records in the SDI database. The
15 next section describes the interface that allows user agents to submit data and update
16 information.

17 The data repository has the following key features: *price rules* are associated with
18 information, to control access to information on the basis of information requested and
19 properties about the requesting agent; meta-information to specify additional information
20 about data records, for example has the information been certified by a third party.

21 We describe the abstract structure of data records in the data repository. The
22 semantics of the data field and the price rules depend on the type of information. All data
23 records, whatever the information that they represent, contain the following elements:

24 1. Owner ID.

25 The owner ID is a three-tuple (Public Key, SDI identity code, Remote address).

26 The public key is provided by an agent that submits data, and is one half of a
27 public/private key pair in a public key based cryptographic infrastructure. The public
28 key is used to provide authentication of the agent, in case it wants to amend the
29 record in the future. The agent can submit a message signed with its public key to
30 prove its identity. An agent can use a different public key for each alias that it
31 maintains within SDI. The SDI identity code is a unique code, generated by SDI for

of currency, or as we also suggest in one useful variation “Community Dollars” that are SDI-specific dollars which can only be spent with vendors that are registered with SDI, and can also be restricted in various ways. As a special case, the prices allow a simple capability-based security system, because they can be set to \$0 or \$infinity on the basis of certificates owned by an agent that requests information. The price-rules are described in more detail in the next section.

5. Meta-information

The meta-information associated with a data record is expandable, but at present we suggest the following pieces of information: is the data is randomized, and is the information certified? The method of random data perturbation is fully described later in this document. Certification can be provided by third-parties, who can verify for example the age or nationality of an individual represented with an electronic profile.

We provide more information on metainformation below.

3.2 Price Rules

The price rules that can be associated with a data element are described in the next main section, the *Query Execution Module* section. The user associates a rule with data that the system uses to compute the price of a query during query execution. The price is further used within SDI to operate a *data marketplace*. Again, this is described in the next section. The user will eventually pay the *total price for its access to all data elements used to compute the final response to a query*; for example negative responses do not incur a price, it is only data that actively makes it into the process of computing the result of a query that matters.

We describe in the *Query-execution* module section the methods that are used to determine which data elements are used to compute the final response made to an agent, and therefore to compute the sum price for the final query. An agent is not charged for every access to data elements made during intermediate stages of executing a query. For example, if a query requests information about ten data records with high value associated with a particular field, the price of the query is the cumulative price for the access to the ten data records returned in response to the query, and not for all the data records queried in determining the ten records to return.

1 In general, a query of a data record can be part of a larger query session, as
2 discussed in the section on query execution. Within a larger query session, the same data
3 record might be queried a number of times, for example with a large compound query
4 split into a number of steps with continual execution contingent on continued correct
5 responses. E.g. Select records of type A, then from those records select records of type B,
6 then from those records select records of type C... In this case, because price rules might
7 be *non-linear* in the amount of information provided, it is important to provide a price-
8 rule method with information to allow it to track a sequential query.

9 As discussed in more detail in the next section, the system of SDI handles this by
10 generating temporary and anonymous pointers to data records, that are valid only for a
11 single session, and allow a price-rule to track sequential queries. The data record pointer
12 allows the history of queries to be recovered, and a new price to be computed on the basis
13 of *total information* provided. Similarly, this is important when a query might collect a
14 lot of information about a data record and then request the identity of the agent that
15 provided the information to SDI. The cost of revealing an agent's identity (and therefore
16 allowing a secondary market in its information) might well increase with the amount of
17 information that has already been released.

18 Similarly, in some cases the same data record might be accessed on multiple
19 occasions, but independently, such that the environment performing queries does not
20 know that the data record is the same record from before. E.g., select records of type A
21 and extract information with rule B, then select records of type C and extract information
22 with rule D. A single record can be of type A *and* of type C, and therefore be selected for
23 information extraction in both cases. In this case, each independent sequential query (i.e.
24 A, B and C, D) has a set of data record pointers, so that the cost of extracting information
25 is computed independently for any record that has type A and type C.

26 Here is a simple example of a non-linear pricing function, that accounts for
27 cumulative information that has been provided to a requesting agent. Suppose that any
28 sequential query receives a temporary and anonymous pointer to a data record, that is
29 only useful in the current query session. The *data record pointer* allows the price rule to
30 implement a non-linear pricing rule. For example, consider the query 'what is the value

1 of field A, B and C', split into queries 'what is the value of field A', then 'what is the
2 value of field B', then 'what is the value of field C'? The pricing rule might state:

3 \$0.1 for any one of A, B or C

4 \$0.2 for any pair of A, B and C

5 \$10 for all of A, B and C.

6 In this case, by tracking the data record pointer with which a data record is accessed, the
7 price of the first two requests can incur an incremental charge of \$0.1, while the price of
8 a third request can incur an incremental charge of \$9.80.

9 At a per data-record level the price of a query depends on the response made to
10 the query. This is important because an affirmative answer to the question 'are you the
11 president of the united states' carries more information than a negative answer. This has a
12 slightly undesirable side-effect, in that when estimating the price of a query in the pricing
13 module, the *estimated price of a query might leak information about the result, even*
14 *without performing the query*. However, we believe that the benefits of linking price to
15 information content in a query outweigh this potential loss in value of information. Note
16 in particular, that in general the ask prices associated with data records are private
17 information and independently set, and therefore are not very revealing, especially within
18 a competitive market place.

19 3.3 Data Submission/Update Methods

20 In this section we describe the basic methods to register with SDI, submit data, update
21 data records, and access payments collected by SDI as information is queried. We use
22 "agent" to refer to the computer system that interacts with the central SDI data
23 warehouse, submitting data and requesting payment from the eBank. Agents may
24 represent individuals, vendors, or other self-interested parties.

25 Data records can be submitted under multiple aliases by a single agent, to provide
26 an additional level of control and flexibility in managing data submitted to SDI. For
27 example, if an agent represents an individual that is an online consumer, interacting with
28 different types of vendors, then perhaps the agent will use two aliases: one for while the
29 consumer is at work, and one for while the consumer is at home. Alternatively, the agent
30 can maintain a number of identities for different activities, or interests of the user. In the
31 preferred implementation agents, for example web-browser based client agents, can

1 manage an agent's selection of identities as it submits data to the central SDI warehouse
2 local to the user. The ability to submit information under multiple identities protects the
3 ability of an individual to prevent another agent building a complete picture about its
4 preferences and profile, while still allowing that individual to leverage as much of the
5 value associated with its information as possible. The SDI proxy agent, situated on user's
6 client machines, implements this functionality. It is described later in this patent.

7 A standard cryptographic public key /private key infrastructure provides a useful
8 technique to implement a system in which agents can maintain different aliases. As
9 proposed in the work of D. Chaum [Chaum 81; Chaum 85; Chaum91] a public key (PK)
10 and private key (SK) serves a number of purposes. First, the public key acts as an
11 identifier for the alias, a name. Second, the agent can compute a new private key/public
12 key pair when it requires a new alias. The agent keeps the private key secure, and this
13 provides a method to allow the agent to validate its identity, for example by
14 cryptographically signing a message with the private key. The signature can be verified
15 with the public key, and the keys can be selected with enough bits to make falsification a
16 computational impossibility. This infrastructure is outside of the current patent, but
17 standard in the art.

18 As described in the JANUS/LPWA system [BGGMM 97; BGGMM 98; GGMM
19 98], it is also possible to associate a public/private key pair with a pseudonymous e-mail
20 address, to allow information to be pushed to an agent that owns information under its
21 alias. The system is implemented via pseudonymous proxies which a user agent to poll
22 and check for new messages.

23 An alias can also be completely *anonymous*, but in this case the value of the data
24 provided may be less in the internal market place implemented within SDI (in the query
25 execution module). The convention for an anonymous alias is that the user agent
26 continues to provide a public key, and use a private key for validation of its identity.
27 However, in this case the public/private key is one-off and just for this data record, and
28 no return address is provided. An agent that submits information anonymously can still
29 recover payments from the eBank.

30 When a data record is first created SDI returns a data ID, so that the agent that
31 submits information can specify a particular data records in future, in case it creates a

1 number of data entries in the central SDI data warehouse. Submitting a new data record
2 to the SDI data warehouse is accomplished with the following semantics:

3 (SDI identity code, SDI record ID) = SUBMIT(alias, data object, price rule,
4 metainformation).

5 The following protocol is followed in the SDI data warehouse in response to a
6 SUBMIT message:

7 1. SDI first checks the alias against its record of existing aliases. If the alias exists, then
8 SDI first verifies that the alias is not anonymous (in which case it should only have one
9 record), and then challenges the agent to sign a random message with its private key—to
10 validate its identity. Once validated, the SDI identity code is returned to the agent.
11 Otherwise, if the alias is not found, then a new SDI identity code is created.
12 2. SDI then creates a new data record, with the data object, price rule, and
13 metainformation provided by the agent, and then computes a new record ID, which is
14 also returned to the agent as proof that the record has been created. This record ID is used
15 to change the data in the future.

16 The owner of a data record can change the record with the following rules:

17 Ok = CHANGE(alias, record ID, data change)

18 Ok = CHANGE(alias, record ID, price change)

19 Ok = CHANGE(alias, record ID, metainformation change)

20 As with the SUBMIT command, first the SDI data warehouse checks that the alias
21 exists, and challenges the agent that submits the CHANGE request to sign a random
22 message with its private key, to validate its identity. Then the record ID is located, and
23 again it is verified that the record is owned by the agent with the alias. Finally, changes
24 are made to either the data record, the price rules, or the metainformation. If everything
25 checks out, then SDI returns TRUE, otherwise SDI returns FALSE.

26 Finally, an owner of data can remove a data record with the following command:

27 Ok = REMOVE(alias, record ID)

28 The checks on the identity of the agent that requests that a data record be deleted
29 are made as for the CHANGE command.

An agent can associate meta-information with a data record, that can serve a number of different purposes. Data records can be tagged with meta-information, that can include but is not limited to, Has this information been randomized? Certificates relating to the data record. For example, the meta-information might be a certificate from a third party about the integrity of information, or a certificate from an agent's client-side SDI proxy agent that the information in the record is unique, and not submitted under any other aliases by that agent.

A key example of the role of a certificate is described later in the patent, where we explain how the client-side SDI proxy can provide a certificate to state that this is the only data record with information X. This is useful, because it allows an agent that submits information to maintain multiple records, but still provide a guarantee to querying agents that certain valuable information is not duplicated across multiple records.

The meta-information may be associated with particular fields in the data associated with a data element, for example specifying that a particular piece of information has been randomized, or that a particular piece of information is highly sensitive and should be randomized before release.

4. Query Execution Module

The query execution module is a key component of the SDI system. Its key functionality is:

- a. Implement an internal market for information as queries are executed, ensuring that information that is provided for a lower ask price is used in preference to information at a higher ask price.
- b. Compute the price of queries based on information finally provided in response to a query, by keeping track of the data records that are used to compute a result.
- c. Interface with the pricing module to allow the price of a query to be computed without executing the complete query.
- d. Interface with the data repository, performing queries either on fields with fixed index labels (i.e. within a traditional database framework), or with respect to an

1 object-oriented framework with queries performed by invoking methods that are
2 associated with data.

3 e. Monitor conditions for the persistent queries, and provide a response to a query if
4 conditions are satisfied.

A central part of the invention of SDI is the method to compute the price of a query as it is executed. We described the semantics of price rules, which price access to data. They are defined by the owner of information. The SDI query execution module implements an internal market for queries, and ensures that queries are executed at minimal cost to agents. The SDI query execution module is also responsible for collecting revenue on behalf of owners of information.

We describe this methodology in this section, describing how the total price of a query is computed, as a sum of the price charged by the owners of data records that provide information which contributes to the final response.

14 The query execution module follows a protocol to execute queries:

- ```

15 1. Request_For_Price(Query, Agent_Certificates)
16 2. Estimate Price with call to the Pricing Module
17 3. Price_Quote(Quote_ID)
18 4. Request_Query(Quote_ID)
19 5. Get_Payment()
20 6. Execute the Query
21 7. Report_Answer()

```

22 In step (1) an agent makes a request for a query to be priced, stating the query, and  
23 providing certificates to allow the query to be priced. In step (2) the query execution  
24 module makes a call to the pricing module, and a price for the query is computed  
25 (described in the next section). In step (3) the agent receives a price quote, and can then  
26 decide whether or not to execute the query, and also a query ID. In step (4) the query  
27 execution module receives a request to perform the query, and then in step (5) requests  
28 payment from the agent. When payment is received the query is executed (6), and  
29 appropriate payment is credited to agents that provide information, scaled to make the  
30 budget balance as necessary (in the case of an over or under price quote). Finally, the  
31 response to the query is provided (7).

#### 4.1 Price Rule Semantics

Each data record has an associated price rule that controls the price of accessing information. A data record can contain a number of different pieces of information, and data can be provided to different degrees of accuracy, so the price rule can be quite complex in general. In particular, consider a data record that represents the profile of an individual. The individual might be happy to have information released about some fields, for example its ZIP code, or its recent book purchases, but less happy to have information released about different fields, for example its salary or social security number. Clearly, different pieces of information within a single data record require different prices.

Furthermore, a price rule for the data record as a whole might need to be superadditive across data elements, such that it becomes very expensive to request too much information about data associated with the same user. One of the driving concerns behind the present invention is that at present it is possible for on-line vendors to collect information about a single individual via “cookies” (identifying codes which are left on a user’s client machine), and form a portfolio of information about various activities and preferences of that individual.

While information about an individual might be acceptable in small amounts, in large amounts the same information can soon become unacceptable. Furthermore, even if users release information anonymously in response to a query, if a lot of information is released the identity of the user can be compromised. Every additional piece of independent information that I state about my profile identifies myself a little more clearly, and acts to distinguish me from the profiles of other individuals. We discuss this further in a later section on random data perturbation, which describes how random noise can be added to data to counteract this effect.

The basic idea is that the system of SDI allows a user to associate a price rule with every data record, that computes the price that a user must pay to execute a query over that data. We allow the price to vary, depending on properties about the requesting agent, the amount of information requested, and the level of identification that is required of the owner of the information in responding to the query.

The abstract form of a price rule for a data record is a function:

6           However, as we describe in the next section, an agent only actually pays this price  
7 if the information provided is used to compute information in the final response provided  
8 to a query. For example, consider query “Select all data records close to record X”. The  
9 query is most simply executed by computing the distance between each record and record  
10 X. Although a negative response from a record that is not close to X provides some  
11 information about that record, the information is not used to compute the information (i.e.  
12 set of records) that is finally provided in response to the query.

As in the representation of heterogeneous price rules in SDI, the invention of SDI allows heterogeneous methods to compute prices for queries. In the most general case, we allow a price rule to be computed as a price method, which is invoked for a data record before a query is to be performed. The query-execution and pricing modules simply invoke the price method, and the data object provided by an agent returns a price for the query. In this most general version, the role of SDI is limited to providing the price method with the query, the certificates of the requesting agent, and the level-of-identification that is required.

In a simpler variation the system of SDI can provide a number of default price rule languages, which allow a user to specify in simple but quite flexible terms a price schedule for queries. We describe two such rule languages: an uncertainty-based additive price rule; and a grouped additive price rule which is a simplified version of the general uncertainty-based rule.

The uncertainty-based additive price rule allows a user to adjust the price for a query based on information about the requesting agent, the amount of information requested, and whether the information is required anonymously or with a revealed identity. It does not allow a non-linear coupling across the price of information about multiple attributes, but instead includes a simple upper-bound on the amount of information that can be requested. Such a non-linear coupling could be added with simple









- 1 • Low price, if an agent presents a Charge certificate, and requests information with
- 2 a High degree of uncertainty.
- 3 • High price, if an agent presents a Charge certificate, and requests information
- 4 with a Low degree of uncertainty.
- 5 • Infinite” price, if an agent cannot present a Charge certificate.

6 Different price functions can be defined for anonymous and revealed-identity. The  
7 price rule has the following components:

- 8 a) A set of certificates that allow data access for free.
- 9 b) A set of certificates that allow data access for some charge.
- 10 c) A list of price-functions, each defined with two prices for revealed-identity and
- 11 two prices for anonymous identity. The pairs of prices are for accurate
- 12 information, and approximate information (with a minimal level of
- 13 approximation).
- 14 d) A mapping of data elements to price-functions.

15 Again, the price of a full query is computed as the sum price over all information  
16 requests to all data elements.

#### 17 **4.1.3 Superadditive Price Rules**

18 We can allow a non-linear interaction between the price of queries on individual data  
19 elements with the introduction of additive interaction prices, which are fired when a  
20 threshold is exceeded for the total amount of information released over all data elements  
21 accessed in the same query. This is a simple approximation to a more general  
22 combinatorial price rule, which would price all combinations of data access terms  
23 explicitly. For example, a rule could state: “If more than T data elements are accessed in  
24 set Sensitive then add a “bundle price” L to the total sum price of the query.”; or: “If the  
25 total weighted access to data elements in set Sensitive exceeds threshold T<sub>1</sub>, then add  
26 bundle price L<sub>1</sub> to the total price of the query”; where the weighted-access is computed  
27 with weights equal to the degree of uncertainty associated with queries on elements.

#### 28 **4.2 Computing the Price of a Query**

29 A query can be quite complex, involving a number of operations on data stored in  
30 the data repository. Call a query session a sequence of operations performed in response  
31 to a query, and before the answer is returned to the agent that submits the query.





0969098-102700

1 separation is maintained between the data and the state of the program that executes  
2 queries.

3 Each query session is associated with a table of data record pointers. The table  
4 records valid data record pointers, that can be referenced in later query operations, and  
5 the current price of queries performed with the data record:

| 6  | Data Record Pointer |  | Current price |
|----|---------------------|--|---------------|
| 7  | -----               |  | -----         |
| 8  | P_1                 |  | \$0.10        |
| 9  | P_2                 |  | \$0.20        |
| 10 | ...                 |  | ...           |

11 The SDI query execution module performs another key book-keeping service, recording a  
12 list of data record pointers that have been used to compute a result in the program that  
13 executes a query.

14 For example, suppose that set PS contains a list of data record pointers, and a  
15 subsequent query operation makes the following command:

16 `_count = Count(PS, 'property X')`

17 which can be interpreted as “count the number of records in set PS with property X”. The  
18 query-execution module maintains a table of variable accounts, which records the data  
19 records that were used to compute a result:

| 20 | Variable Name |  | Associated data record pointers |
|----|---------------|--|---------------------------------|
| 21 | -----         |  | -----                           |
| 22 | _count        |  | P_1, P_4, P_6                   |
| 23 | _tmp          |  | P_3, P_8, P_10                  |
| 24 | ...           |  | ...                             |

25 Whenever the value of a variable is computed from the values of a number of existing  
26 variables, then the list of data record pointers is the union over the data record pointers  
27 which were used to compute the values of the other variables.

28 Eventually, when a result is reported at the end of the execution of a query, the  
29 price of the query is the sum of the requested payment for the data record pointers that  
30 were used to compute the information provided, as stored in the table of data record  
31 pointers.

Figure 3 illustrates the method diagrammatically. We maintain a “firewall” between data and the state of the query execution, with all access to the values in data records via data record pointers, and invoking price rules associated with data records. On the query-program side we maintain a table of data record pointers, and the cost of the cumulative information received from that data record under that pointer, and a table of variables and the data record pointers which were used to compute the values in the variables.

All manipulation with the values of variables on the “query-side” of the data firewall can be performed with no additional data-access charge, because data access charges are only incurred in computing values of variables, when that directly or indirectly requires that additional information be collected about the data records.

The only charges are incurred when the processor receives information. This information might be in the form of a number, that represents the count of records with a particular property, aggregate information in the form of a new collaborative filter, or raw data extracted from data records.

For example, SELECT operations are free until information in the data record is used to compute the value of something which is used in computing the final response returned in answering a query. A charge is accounted for against the data records which are queried and selected, but that charge is only levied if information is used.

#### **4.1.1 Internal Market within the Query Execution Module**

The data side of the firewall is also responsible for performing the cheapest queries possible. For example, if there is a request to retrieve 100 data records with property X, then the database retrieves the data as cheaply as possible based on the access-control rules of the owners of the information in the database.

All operations are implemented within an internal market. The internal market ensures that information is provided according to the price rules associated with data records, and also places providers of data in competition with each other.

The system of SDI implements an internal market, with queries executed as cheaply as possible. Both the SELECT and the PROCESS commands allow an associated restriction on size, e.g. “SELECT 100 records with property X”, or “PROCESS 10 of the records with data pointers in list X1 with method Y”.





1 In step 2 the data records that relate to profiles of users that have recently  
2 purchased a flight are now sorted in terms of their salary field. An algorithm to sort  
3 performs a sequence of comparisons between data records, with each comparison  
4 requiring that more information is revealed about the underlying salary of the profile, and  
5 changing the price associated with the data record. The prices are updated in the data  
6 pointer table according to the data record price rules, and the query program now has a  
7 list of ordered data record pointers.

8 Now, in step 3, the query program initializes a new set of data record pointers.  
9 This includes all the pointers for users that are willing to have their identities revealed.  
10 The pointers are maintained in order of salary. At this stage the entry for this new  
11 variable in the table of variable accounts only associates the data records with this new  
12 property.

13 Finally, in step 4, a new variable is initialized to contain the first 100 data record  
14 pointers from step 3. These are the data record pointers listed in the variable accounts  
15 table. These are the only data record pointers that receive payment as a result of this  
16 query, because the intermediate information provided by other data records remains  
17 within the SDI query execution module and is not released to the agent that performs the  
18 query.

#### 19 4.3 Discussion

20 The pricing model as outlined above has the following attractive property:

21 The price of a query is independent of the order in which operations are  
22 performed, and optimization of order is not necessary for reasons of price.  
23 This is nice, because it means that agents do not have to perform lengthy optimization to  
24 structure queries in cost-efficient ways. The pricing module also ensures that:

25 Agents only pay for information received as the result of executing a query.  
26 Finally, we implement an internal market, so that with all other things equal, data mining  
27 and other types of queries are executed as cheaply as possible:

28 An internal market favors data records with low cost pricing instead of records  
29 with high cost pricing, with all other things equal.

30 As an example, consider this query:

31 COUNT (SELECT 'FIELD1 = X' and 'FIELD 2 = Y')

Similarly, consider the query:

`SELECT(Z, Select(Y, Select(X))); return '2'.`

i.e., perform some complex nested query but return the fixed value '2'. The query execution module tracks that no data pointers are used to compute the value returned to the agent, and the querying agent is not charged for the information provided about data records to compute the `SELECT` operations, because that information is retained within the central SDI data server.

In a SELECT query the agent can only ever pay for information provided about data records that respond positively to a SELECT, because no information is ever released to agents about the data records that respond negatively.

In another example, consider a query ‘COUNT the data records with PROPERTY X’, that returns the result ‘none’, i.e. there are no data records with this property. This query is priced at \$0 because none of the agents that own information provide information to compute the

Therefore, a query of the type ‘how many data records have type X’ that answers ‘none’ has price \$0, because the individual agents have not revealed any information. Similarly, because we would like all equivalent queries to have the same price, a query “how many data records do not have type X” in the same database (answer “all of them”) should cost \$0 because the same query could have been formulated the other way. A reasonable way to ensure that a query is formulated in the cheapest possible way is to use an query optimization manipulation program, of the type that is common in database research, during the ‘price quote’ phase in order to find the formulation with the cheapest price.

The pricing module is a core component of the central SDI data warehouse. It has two functions: (a) to *predict the price of a query without executing a complete query*; (b) to aid in the allocation of payment given a price quote.

We build into the pricing module a method to allow an agent to make a quality/cost tradeoff before executing a query, and even to specify a restricted query that will respect an agent’s budget. Already embedded in the query-execution module is an internal market for data records, to ensure that cheaper data is used in preference to more expensive data access, given two identical data records. Furthermore, already embedded in the query-execution module is that agents only pay for the *provision of information*. Finally, the method computes a price for a query which is independent of the order of operations.

We propose the following method for providing initial price quotes to agents in SDI:

1. Compute a range of price quotes in response to a query, for different amounts of computation (which can be varied by limiting the input set of data records), and for corresponding values of a “quality metric” provided by an agent that submits a query.
2. Allow the querying agent to select an optimal price/quality/computation tradeoff, and then receive payment, execute the query, and provide the response.
3. Allocate payments to accounts in the eBank that correspond with data records which are used to compute the value of the result of the query, with payments adjusted if necessary to account for any error in the price quote.

### 5.1 Computing a Price for a Query

We allow a small *fixed price*  $F$  for an initial price quote, to cover the cost of computation and prevent agents saturating the server. The price module computes a *price quote* based on a statistical technique to compute the price of the full query without executing the full query.

The basic idea is to perform the query on a number of records, e.g. 1%, 5%, and 10% of the total records that will be used for the full query, and then extrapolate to the full query size. We need to perform the complete query because on a subset of data records, instead of part of the query on all data records, because query programs can have different phases—and each phase might incur very different information costs. The



number of records that are used for the query can be restricted by making a random sample of the total domain of data records used for a full query.

## 5.2 Allowing a Price/Accuracy Tradeoff

In computing the initial price quote we also suggest a tradeoff for the user between “result quality” and cost. *Quality is subjective, and usually best measured by the agent that requests a query on the data. Therefore, we suggest that the querying agent should provide the metric, but limit this metric to a scalar value to prevent information leakage.* We allow an agent in formulating a query to specify a key metric that it will use to choose an appropriate size of query to execute.

The core technique that we use to control the accuracy of the result of a computation is to limit the number of total data records that a query runs over as input to the query. Take a random selection of all the data, and use that for the processing. At this initial step we limit the price and the amount of computation performed.

As an example, consider a query to collect pseudonyms for agents with useful properties. The metric of interest in this case is the average per-pseudonym cost, and this can be provided by the pricing module *without providing any useful information to the querying agent*. Given this information the agent can then decide whether to proceed with a query, and how much money to allocate.

Similarly, perhaps a querying agent cares about the quality of match between data records and data records with ideal properties. The agent can provide a method to *instrument* its query to compute this quality, and the price module can in pricing the query provide guidance about the tradeoff between running the query over different numbers of data records.

Another more involved method is to introduce randomization as the query is processed, because some owners of information may provide randomized information more cheaply than exact information.

That will vary depending on the number of data elements that are queried.

### 5.3 Making a Price Quote for a Budget

Consider the problem of an agent that wished to perform a query with a budget  $B$ . For example, I would like to spend \$1000 to compute a collaborative filtering model. There is a particularly simple way to allow this:

First, estimate the cost of building a collaborative filtering model on the basis of  $N$  data records, selected at random (or with some preprocessing of the database to identify useful data records). Estimate price  $P$  for the query.

Second, adjust the number of data records used to compute the actual query based on the price  $P$  that was computed, and the budget  $B$ . The number of data records, assuming a linear model of per-data record charging is simply  $N \times B/P$ .

If there is reason to believe that the cost of a query is a non-linear function of the number of data records, then a number of samples could be priced, for suitable sample sizes  $N_1, N_2, \dots$  etc.

Ultimately, the result is a query that is adjusted in scope by an upfront limit on the number of records used for the query domain. The price module can scale the cost of the query, and make an appropriate price quote.

#### 5.4 Optimizing the Price of a Query

We also suggest a technique to adjust the formulation of properties in a query, e.g. 'COUNT(SELECT records with property X)', could be reformulated as 'Count(all records) - COUNT(SELECT records without property X)', possibly for a smaller price. A random search technique can provide one simple approach to optimize query formulas.

## 6. Communication Module

The communication module in the central data warehouse takes incoming messages from agents and decides how to handle them, and sends responses to queries back to agents.

The action selected in the communication module depends on the *type* of the message, for example the message can be a data-update request, a new query to execute, a request to add new profile information, a request to find agent profiles that match, or a request for processing a set of profiles that meet a particular criteria with an application in the database.







- 1       \* Pagers
- 2       \* Wrist watches (typically one-way, receive only)
- 3       \* ATM machines, point of sale kiosks
- 4       \* Ceiling-mounted video cameras
- 5       \* room-based motion and heat detectors
- 6       \* Internet-enabled automobile
- 7       \* Chip-enabled appliances
- 8       \* Electronic Schedule Books

## Smart Home Application

The concept of a “smart home” is not a new one: Even in the 1950’s futurists discussed the possibility of imbedding intelligent systems into a house, relieving the homeowners of such boring daily tasks as watering plants or adjusting the thermostat. Such a system could easily be implemented today using control computers running simple sets of IF-THEN rules, but might soon prove unreliable because of its inability to adapt to new situations. For example, a garden party could easily be ruined if inadvertently scheduled at the same time as a scheduled lawn-watering.

LEIA provides a much more sophisticated alternative. Given a house that incorporates an internal local area network (LAN) linking household appliances, sensors, communication devices, and home computers, LEIA could easily form the basis for a truly intelligent and adaptive home environment capable of handling daily chores, security, and information delivery.

Firstly, LEIA could easily subsume the tasks traditionally slated for use in a smart home; however, its ability to infer the location and needs of family members would greatly enhance the handling of these tasks. For example, the above-mentioned garden party would not be spoiled because a motion sensor would inform the watering system that a party is taking place on the lawn. Or, being set in a power-saving mode, LEIA could light only those rooms currently being used; as a family member moved around the



1  
2       iii) appliances

3       \*alarm clock

4       \*light switches

5       \*stove

6       \*dish washer/clothes washer

7       \*water taps

8  
9  
10       In addition, accessed via secure channels (perhaps using a personal agent as an  
11 intermediary), three major types of personal data are also useful for running the smart  
12 home:

13  
14       i) Explicit rules for routine tasks (E.g., desired schedule for lawn watering)

15       ii) Expressed preferences (E.g., a desire for morning news about international  
16 events)

17       iii) Historical database of previous interactions between the family and the house  
18 (This will  
19 allow LEIA to learn daily routines and family habits).

20  
21       There are two major classes of data that need to be served by LEIA: (1) control  
22 signals for household systems and appliances, and (2) information desired by household  
23 members.

24  
25       To a large extent, household systems can be appropriately controlled using a  
26 traditional rule-base, although a fuzzy rule-base would be inherently more flexible (IF  
27 <the bedroom is somewhat chilly> AND <it's almost time to wake up> THEN  
28 <substantially increase the bedroom's heat>). Many of these rules would be explicitly  
29 programmed by the family from the beginning; others could be learned using inductive  
30 methods.





complementarity) to the print media viewed by the user presently. Alternatively, the data captured by the miniature camera device relating to the particular pages of print media presently observed by the user may consist exclusively of an identifier code a portion of which is used by the system for purposes of identification of the code (or alternatively such identification may be determined by its physical coordinates on the page itself). The other portion of the code is a unique identifier of the particular page, which may be used to identify manually by a content expert whereby select print materials and particularly their associated electronic counterparts are either matched, designed and/or customized for use within their system context of the present convergent media application. It is apparent that such a system could also be usefully deployed to benefit advertisers, e.g., presenting advertisements electronically to correspond with particular printed content or matching printed ads with corresponding electronic counterparts.

This complimentary content, be it audio, video, text, or combinations thereof, including advertising conforming to user-defined conditions, is delivered through various in-home devices located in the proximity of the user.

### **Smart Office Application**

Large business centers can be very confusing places: people are often away from their desks, many different meetings are scheduled and rescheduled during a single day, people are called off to emergencies, and clients drop in for unexpected visits. The inability to coordinate people and information in such situations leads to greatly reduced efficiency.

LEIA offers a unified solution in the face of this chaos; by linking together people's schedule calendars, communications systems, contact information, and location coordinates, a system can be devised to (1) shepherd a visitor (even if unexpected) around the business center, (2) push vital contact information about the visitor to those in the company who will be visited, and (3) adaptively coordinate important meetings within the company.

**When a visitor comes to the center, he will:**

b) receive a hand-held computer that, while presenting the visitor with useful information, will emit a constant signal enabling LEIA to pinpoint his position at all times.

On the employee side, LEIA can acquire signals from:

a) active badges

b) telephones

c) beepers

d) interactions with terminals

All sensors are connected directly to a LAN internal to the company. Only basic security measures need to be enacted.

Registered visitors will likely be profiled in standard contact files containing information about their own company, position, professional interests, previous interactions with the company they're visiting, and basic personal details (e.g., photograph, wife's name).

Employees of the company using this LEIA-based system will have much more information available. In addition to the standard contact information, LEIA will have access to their phone and e-mail logs, daily calendar/schedule, work group and project details, position within the company hierarchy, and work-station notes (e.g., an unexpected “out to lunch” message could be posted on-screen by the employee before they run off).

Although some higher-level intelligence is needed (e.g., for matching interest profiles between employees and visitors), the office-based system could well be run with



Given that multiple employees have signed up for a meeting, or that a single employee has sent out a signal indicating the need for a meeting and its priority (standard or emergency), LEIA gets to work coordinating the details. Cell phone and beeper signals give employee locations outside the office, active badges signal employees already at the office. Calendar agents are contacted for availability; a standard meeting might depend on people having time slots free, whereas an emergency meeting would be given priority status, and could bump items already scheduled (given that they have been assigned low priority).

An optimal meeting time and place (or teleconnection for remote employees) is set up, and any pre-meeting information is forwarded to the appropriate workstations.

meeting, engaged in an important phone conversation

### **Resolution Credentials**

Section 2.2.1 of this patent discloses the use of resolution credentials for the facilitation of agent-agent interactions. Resolution credentials enrich the quality of such interactions because they certify the status claimed by individuals, allowing agents to maintain relationships of trust. On a more practical level, resolution credentials are of use for matching and introducing users to each other on the basis of common interests, for guarding access to high-level users (with the agent acting as a gate-keeper), and for flagging times at which an individual is reachable (and not tied up in meetings, or the like, as specified by the scheduling agent).

## **7.2 General Description**

We allow a user of SDI to push some data **A** to the data warehouse, and make this data available to all other agents via the SDI query-execution module, and to retain local control over other data **B** on the client machine. In this section we discuss this and other variations, and explain how the functionality of the central SDI data warehouse can be

replicated on a client machine in a *client-side SDI database*, to allow a user to leverage the value in local data when interacting one-on-one with another agent. For example, a consumer’s shopping agent might wish to use local data about its preferences to allow an online vendor to configure its services, but without releasing the information to the vendor. We explain how this can be supported within SDI.

We allow the following variations:

1. A user can store some information on his/her local client machine, with that information provided to other agents directly, or retained for querying by other agents on the local client machine as in the system of query-execution on the central SDI database.
2. A user can use a light client machine, perhaps a mobile computing device, which is periodically configured with information stored on the central SDI database that is relevant to the current task. This information can then be used directly by other agents in interactions.
3. A client machine can be automatically configured with information about a user when the user is in its proximity, again to allow direct information exchange with a local agent.
4. The client machine can maintain all information, with virtual links from the central SDI data warehouse to the information. The system operates as before, with information provided if conditions are met, etc.

There are two basic motivations at work. One, a user might prefer to keep control over some information, even keep control away from a central SDI data warehouse, in case it is compromised in some way. We can allow a user to maintain data on a local client machine but seamlessly receive the same benefits as if the data was pushed to the central SDI data warehouse, with virtual links to the client machine. Query-execution can

operate much as before, with data pulled from the client machine as necessary and charges levied. Operating overhead is incurred in such a system because data must be communicated over the Network during query execution. The client machine must also be “fat”, with enough local data storage.

A different motivation occurs when a user has a very light client, for example a mobile computing device such as a Palm computing device or a cell phone. In this variation the user will perhaps like to store profile information relevant to the profile that it assumes with another agent in a one-to-one interaction, but not have space to store all profiles. The role of the central SDI data warehouse is to configure the light client dynamically.

With information on a user's client machine, then we can support local queries from other agents, for example on-line vendors might request information to allow an appropriate product to be displayed to a user, or a relevant advert. The queries can be executed in a number of interesting ways. First, the information can simply be released to the agent, and the agent can make appropriate use of the information. Second, the information can be retained on the client machine, and the agent can submit its *selection function*, and receive the optimal selection, for example a good advert to show the user.

One interesting implementation of such a client-side data warehouse is via *smart stick technology*, developed by Sony, which is a robust “data wand” that allows the storage of configurable data.

### 7.3 Important Data Location Variations

It is useful to identify four key modes of interaction between a pair of agents within the system of Secure Data Interchange. The agents might, for example, represent a consumer and a vendor in a B2C e-commerce application. The modes of interaction are adjusted to allow for different data locations.

### 7.3.1 Data in SDI data warehouse

In Figure 5 we illustrate the situation where agent 1 has its personal information about the user stored in the central SDI data warehouse, and provides agent 2 with the ability to process a query on the data warehouse and access information about the agent. The query might either request information about the agent’s profile, or provide a method to execute in the data warehouse, where the result of the method is an optimal action for Agent 2 to make in providing personalized information and/or services to the user with agent 1.

### 7.3.2 Data on the Client Machine (client-side SDI database) and in SDI data warehouse

Figure 6 shows an alternative variation where data is stored on an agent's client machine and in the SDI data warehouse, and the client-side SDI data warehouse can respond to queries from agent 2 in the same way as the central SDI data warehouse. In particular, as we suggest later in this document, this variation allows an agent representing an on-line consumer to store very personal information on a local machine, and for example only allow other agents to benefit from the *results* of analysis on that information without receiving the actual information.

### 7.3.3 All information stored Client-side.

Another variation on the basic SDI architecture, is that all information is stored only as *virtual links* within the SDI data warehouse, with physical storage on an agent's local client machine. Preferably, in order to leverage the value of information the client machine would be always networked with the central SDI data warehouse. A current technical solution which seems appropriate would use a *Set Top Box*, connected via cable TV lines to the SDI network. Many households in the US have set top boxes, and this



1 would provide a quite convenient way to leverage the value of information but keep it  
2 secure against inappropriate access in a central data warehouse.

3  
4 In this variation, to support a personalized session with another agent the data is stored  
5 locally, and can be processed using a client-side SDI data warehouse with a local query-  
6 execution module. To also support data mining by agents of a shared database, the client-  
7 side data records could useful push *data types* and *price rules* to the virtual shared  
8 database, to allow more efficient searching. This variation describes a *distributed SDI*  
9 *data warehouse*, with indexing performed over on-line client machines.

## 11 8. Distributed Query Execution

12  
13 In addition to allowing distributed data location, the system of SDI allows *distributed*  
14 *query execution*, via client-side query-execution modules. The advantage with this  
15 variation is that data can be physically stored on trusted local machines close to a user of  
16 SDI, and the queries can be physically executed on the machines so that not even  
17 intermediate results are available outside of a local firewall protected system. There is no  
18 technical solution provided within SDI to prevent agents sharing information to other  
19 agents, other than to prevent agents from receiving that information in the first place. We  
20 do this using three many techniques: pseudonyms, information randomization (see  
21 Section 10), and query-execution in the SDI central server without release of raw data  
22 inputs.

23  
24 We can implement safe client-side query execution with direct replication of the key  
25 functionality of the query-execution and price-module of the central SDI data warehouse  
26 on a user's client machine. A querying agent can push the same method to an agent's  
27 client machine as it would submit to the central SDI query execution module, and an  
28 agent can execute the method locally and release information consistent with his/her data  
29 release rules. The SDI client-side agent can be configured, via the profile management  
30 and data-release policies, to control the type and amount of profile information released

The challenge is to be careful that the results of a query do not reveal too much information, but this is solved using the same *price rules* as in the central SDI data warehouse. The mode of interaction supported is QUERY(PID), i.e. execute a query on a particular pseudonym ID, and the client side query execution module ensures that the vendor has a certificate from the client-side SDI agent to query the information in the local database about that pseudonym. The control over the profiles that a vendor can access, coupled with the price rules to ensure that information is perturbed as necessary and that not too much information is released, and to allow an agent to leverage the value of information.

A user's agent can retain control over local information by limiting the information that can be released in response to a query. For example, although a query can be complex and access a lot of profile information on a user's client machine, we suggest an additional protection that controls the information that a vendor can receive for future use. It is possible to limit the response that the vendor receives to a fixed message size, e.g. to 10 bits for example, to provide a very strong overall control on the ability of a vendor to use information again in the future.

### 8.1 Example: Customized Vendor Web Pages

The method, called *safe client-side query execution*, has a direct application to systems for personalized information delivery, where it is not desirable to release profile information to an information provider, but preferable to allow the information provider to provide methods (e.g. queries) that are executed on a client machine, and personalize information before it is provided to the end user.

A key application of this technique is in B2C on-line e-commerce where a vendor can push methods for personalization of a virtual shop front to a user, and the user's SDI enabled client machine can implement the methods (which are queries) on the client, and

1 push the result back to the vendor's server. The server then constructs and pushes an  
2 optimized virtual shop front to be displayed on a user's local display.

3  
4 At present advertising networks such as DoubleClick ([www.doubleclick.com](http://www.doubleclick.com)) are able to  
5 track a user as he/she browses across multiple sites in the DoubleClick network because  
6 the ad server can place a "cookie" text string on the hard drive of a user's client computer  
7 which identifies that user as he/she browses. SDI is designed to be used on-line in  
8 combination with a method to block cookies, replacing them for example with "safe  
9 cookies" [Netscape 96] to still permit stateful interactions during a single session with a  
10 vendor, for example "shopping basket" style interfaces.

11  
12 The SDI-client side proxy agent can replace the role that cookies take in automatic user  
13 log-in and password checking, through a Janus/LPWA-style implementation [BGGMM  
14 97; BGGMM98; GGMM98]. A user's log-in user name and password can be derived  
15 from a user's pseudonymous identity, and computed within a cryptographic framework.  
16 Janus also provides a technique to allow a user to receive e-mail pseudonymously.

17  
18 A core technique in SDI is to use a unique pseudonym for each vendor, and then  
19 selectively provide vendors with profile information from across multiple pseudonyms;  
20 either anonymously within the central SDI data warehouse for data mining purposes, or  
21 alternatively via client-side personalization so that a vendor can leverage a user's wide  
22 profile data without receiving *direct* access to that information.

23  
24 In providing profile information to a vendor during an interaction, so that the vendor can  
25 provide targeted products and servers to the user based on data mining that it has  
26 performed on aggregated user data in the central data warehouse, the user's agent must be  
27 careful to protect the identity of a user. This means that the agent must not reveal  
28 information under pseudonym P\_1 and information under pseudonym P\_2 to another  
29 vendor that allows the vendors to link the pseudonyms and reason that the agent  
30 represents the same user.

1 We describe in Section 10 the technique of random noise perturbation to release  
2 information in response to queries, and another simple technique is to carefully protect  
3 particularly sensitive information (e.g. social security numbers, etc.)  
4

## 5 8.2 Client-side Query Execution 6

7 As described earlier, a key variation of SDI retains a local data warehouse and query  
8 execution module that contains information specific to a single agent, on the client  
9 machine of that agent. It is then possible, via the same query execution controls as  
10 described in the central SDI data warehouse, to allow vendors to characterize and profile  
11 a user based on its local data, but without gaining explicit information about the user's  
12 local data. The vendor can push appropriate methods and targeted services. We describe a  
13 number of bottom-level applications later in the patent, including personalized web  
14 pages, ad-networks, etc.  
15

16 We describe one variation in which the vendor provides generic information to the client-  
17 side device, that filters that information locally. This is relevant for example in high  
18 bandwidth information services such as digital television and satellite systems, where a  
19 user's set top box can store program information locally and make local decisions about  
20 what programs a user might like to watch and when, without providing any profile  
21 information to the head-end server.  
22

23 In this variation an agent can submit a query to be executed locally on a *client-side SDI*  
24 *data proxy*, that stores information submitted by the local agent. The applications of this  
25 technique are exciting, because it allows a user to maintain even more control over  
26 profile information but still benefit from personalized interactions with vendors. The  
27 client-side SDI data proxy just contains data for the local agent, and allows agents to  
28 submit 'personalization queries' of the type 'execute this profiling rule and tell me what  
29 product to show the user'. The outcome of the processing might be personalized  
30 information, for example a selection of books that a user might be particularly interested

1 in based on its profile information and the methods passed to the agent by the agent that  
2 sends general information about the books in its catalogue.

3  
4 Figure 7 illustrates the process. The SDI proxy agent maintains a set of profiles for a user,  
5 relating to transactions that the agent has performed with other agents in the system, and  
6 also other information that relates to the user associated with the agent. A vendor agent  
7 can send generic information and a method for personalization, which is executed as a  
8 query on the local profile data base, with the results allowing a vendor to decide how to  
9 target its products and services to a user. For example, a book retailer might be able to  
10 provide a selection of books that a user might be particularly interested in based on its  
11 profile information and the methods passed to the agent by the agent that sends general  
12 information about the books in its catalogue.

13  
14  
15 We can provide *rule templates*, which are used to select the type of method that the  
16 provider of the information is requesting be applied to the generic information that is  
17 passed to the requestor agent. The provider agent only needs to provide *parameters* for  
18 the processing. For example, if collaborative filtering is the desired personalization  
19 technique then the providing agent provides a list of prototypical cluster centers, but  
20 does not need to specify a collaborative filtering algorithm. The interpreter takes the  
21 information and methods from the providing agent and selects the appropriate rule  
22 template to form a program that is executed as a *safe query* in the CPU, with the usual  
23 controls over the amount of data that a query can access.

24  
25 The output from this processing is personalized information, that can then be displayed to  
26 the user, perhaps to a final test in the control module for applicability.

### 27 28 **8.2.1 Preventing Information Leakage**

29  
30 The challenge is to be careful not to allow the *results* of a query, even those data is not  
31 explicitly revealed to an agent that submits a query, to carry sensitive information that a









## 9. Data Management

The invention of SDI, in its most general form, refers to the system for a privacy-protected market for data interchange between multiple parties. In this section we discuss possible methods for an agent to manage its disclosure of data to the data interchange, and to other agents in the system. We allow agents to adopt different identities, and submit different information under different identities. The idea is that a pseudonym prevents other agents from exchanging information. Given an identity management policy, the *data management policy* is then used to determine what data to submit to the central data server, or release to other agents, under a particular pseudonym. The agent needs to choose *what* data to release, and what *price rules* to associate with that data.

We propose client-side “proxy agents” to intermediate the interactions between agents, and manage data provided to the central SDI data warehouse; and methods to manage an agent’s identity in interactions with other agents, for example via pseudonymous and anonymous interactions. The proxy agents control the amount and types of information exchanged between agents. We refer to the proxy as “client-side”, because it resides on the machine local to the agents that participate in the system of SDI. The client-level SDI proxy, implemented as a client program running on the user’s client machine, manages all data release from the user to other agents and to the central SDI data warehouse. The proxy might usefully provide a rule-based interface to allow a user to select appropriate data management policies.

An important application of SDI is to a system for business-to-consumer (B2C) e-commerce, where SDI allows individuals to provide vendors with access to profile information that is collected client-side about an individual as he/she interacts with vendors, but retain control over the amount and level of detail that is made available and collect payment for that information. A key function of the client-level proxy is to control the ability of multiple vendors to track a user, by control of the identity of a user in its interactions and the information released. The client-side proxy agent maintains profile information about a user, but releases that information according to a user-defined policy.



agents to match its identity based on information that it releases, then the other agents have no way of combining information about the agent. In the business-to-consumer e-commerce example, vendor A cannot share information about consumer Z with vendor B unless vendor B has some way of linking its customer with the customer of vendor A. This can be done at present via cookie mechanisms and advertising networks (e.g. DoubleClick), or simply by matching identifying information (e.g. user name, e-mail address, credit card number) across sites.

In Section 9.3 we describe data-release policies, which are methods that determine what information is released for each pseudonym a user selects, and the rules associated with accessing that information. We also note that an agent can choose to introduce a level of noise-perturbation to data before its release, so that the data is still useful but cannot be used to link an agent across multiple pseudonyms via the data that it releases.

Alternatively, an agent can release data with a reduced level of accuracy, for example within a range of values, and/or restrict the types of data that are released. Methods of data perturbation are discussed in more detail in Section 10.

Finally, note that there are two parts to the data management policy. One relates to the data that a user releases to the central SDI data warehouse under a particular pseudonym, and another relates to the data that an agent releases to another agent in a direct agent-agent interaction. The following bullets summarize the distinction, and note the sections for more details.

- [9.3: Data Management Policy] Manage the Release of Data to the Central SDI data warehouse (with either physical uploading of data, or release of links to physically distributed data); this includes the association of price-rules with data that is provided to SDI.
- [9.4: Agent-Agent Data Management] Manage direct Interactions with other Agents, for example controlling the identities used in interactions, and the amount



1 system [BGGMM 97; BGGMM98] allows a client-side proxy to compute a new pseudonym  
2 for an agent, and also a new e-mail address and user-name and password for that  
3 pseudonym, to allow access to web-pages that require user log-in. In particular, it is possible  
4 to compute new pseudonyms for a user without a central register of pseudonyms that could  
5 compromise a user's identity, through the technique of "blinded signatures" [Chaum 85].

6  
7 The identity-management policy can be configured by a user when he/she first registers with  
8 the system of SDI; for example a simple policy might classify agents according to the  
9 certificates that they can present, and select a pseudonym according to the agent's class.

10 Other reasonable identity-management policies include:

- 11  
12 • Absolute Privacy. At the highest level of privacy an agent interacts anonymously  
13 with every other agent so that agents cannot learn about the agent across sessions,  
14 and cannot personalize future interactions. Furthermore, information about the  
15 transaction with the agent, cannot violate a user's privacy when provided to other  
16 agents, so long as the agent does not release any identifying information during its  
17 transaction.
- 18  
19 • High Privacy. At the next level of privacy, an agent interacts with every agent under  
20 a unique persistent pseudonymous identity. This prevents another agent exchanging  
21 information about the agent with other agents, unless the agent reveals other  
22 identifying information. However, this policy does allow agents to receive  
23 personalized information over multiple sessions with the same agent, directly  
24 without the vendor using information submitted to the central SDI data warehouse,  
25 or stored on a client-side single-user database. Another agent can personalize  
26 information, but only on the basis of previous transactions with the agent. In an  
27 Internet environment this mode of interaction allows a vendor to track its customer  
28 across multiple sessions.
- 29  
30 • Medium Privacy. Use the same persistent pseudonymous profile with groups of  
31 other agents, perhaps segmented according to the type of task that an agent is

performing. With this level of privacy agents still control of the ability of other agents to build portfolios of information. Examples: (a) an agent might specify one identity for whenever it is purchasing books, so that any book seller agent that it approaches can access profile information that relates to the agent's previous book purchases and queries (so long as someone will provide/exchange/sell/rent that information to the agent); (b) an agent might require that a different pseudonym is used for business-related transactions and leisure-related transactions, to allow more precise personalization; (c) consider a family that interacts with other agents through the same SDI proxy. Each member of the family might choose to use a different pseudonymous persistent identity for all of his/her on-line transactions.

- Low Privacy. No attempt is made to protect the identity of an agent. This is the default privacy level provided via current Internet browsers, at least in the common usage of individuals. The system of Secure Data Interchange is useful with this level of privacy only to the extent of leveraging the value of information that has not been released to any agent. There is not technical method to prevent agents from exchanging information that relates to transactions they have performed with the same agent.

A general method to select a "medium" privacy strategy, with different pseudonyms for different agents, but some shared pseudonyms, is to suppose that agents present certificates. A certificate places an agent into a particular class, and each class is associated with a particular pseudonym. Therefore agents that can position themselves in a particular class that is shared with other agents can share profile information; while an agent that can not position itself within a shared class will receive its own unique pseudonym.

A useful default policy is to suggest that an agent adopts a unique pseudonymous identifier for each vendor that they interaction, and then explicitly link pseudonyms as they decide, via a link-command to the central SDI database. The policy is a hybrid of the core suggestions above. Initially a user selects high privacy, but the user can choose to incrementally relax this privacy decision on the basis of continued interactions with agents. A user might also

sell links between profiles, to agents that can present the right characteristics and pay the right price.

**Example in Consumer-Business On-line Commerce.** Figure 11 illustrates agents 1, 2 and 3: agent 1 maintains one pseudonym for vendors 1 and 2, and a different pseudonym for vendors 3 and 4; agent 2 maintains one pseudonym for vendors 3 and 6; agent 3 maintains a unique pseudonym for each of vendors 1, 3 and 5.

Of course, a vendor might adjust its level of service according to the level of identification provided by an agent. For example, a vendor might provide a more complete service to a user that provides a pseudonym than a user that provides an anonymous identifier, because a pseudonym allows a vendor to collect information over an extended period of time and build a detailed profile that can enable future personalization.

### 9.3 Data-Release Policy

The client-side SDI proxy provides support for automatic submission of profile data, and other types of data, to the SDI central data warehouse. This includes policies for pricing queries, and policies for introducing random perturbations to data.

There are a number of human/computer interactions that are important in allowing a user to configure these options. . Many of the issues were anticipated in Cranor's work [CR 98; RC 99], in which she describes methods adopted in the W3C P3P (platform for privacy preferences) project, and suggests that users can as a first step select a special interest group with which they are affiliated, use that as a base policy, and then adapt the policy as necessary. Alternative techniques include decision-tree methods as described in pending patent [INSERT US PATENT NUMBER AND NAME], that ask a user a number of questions in order to ascertain an appropriate service. There are many dimensions that a user might like to identify; for example: the type of information that

1 can be released, the types of queries that can be performed on that data, the price that  
2 must be paid to perform the queries. For example, we might hard code different price and  
3 data-release policy codes, e.g. A B C D..., and allow third-parties to provide maps  
4 between data types and an appropriate policy code.

5  
6 We view an agent's complete data release policy as an intersection of its data  
7 management and identity management policies. As with identity management, a simple  
8 variation will define a fixed number of data management policies, stating types of data  
9 that can be released and conditions, and a mapping from agent certificates to policies. An  
10 agent can receive whatever policy it presents a certificate for.

### 11 12 **9.3.1 Data Validation**

13  
14 The SDI client-side proxy can also provide a useful service of validating data in profiles.  
15 This can be important in scenarios in which it is possible that a user can try to misstate  
16 information in the database for personal gain. For example, consider a consumer-business  
17 e-commerce system in which vendors determine user discounts for products based on  
18 their profile information. Clearly all users will attempt to adopt the profile for the  
19 cheapest price, if that is known, and if that is possible. We can allow vendors to specify  
20 that they will only follow validated information, for example if a user's profile states that  
21 he/she is very price sensitive, then it is more useful to know that if the assessment is  
22 based on validated transactions performed via the SDI system. In general terms,  
23 information is valuable within a marketplace if the information is accurate, and we can  
24 use client-side validation to achieve that goal.

25  
26 Consider another business-consumer e-commerce example. In general it is useful to allow  
27 an agent to duplicate information across different identities in the SDI data warehouse,  
28 e.g. my Zip code might well be information that I am very happy to associate with all of  
29 my profiles because it is not too revealing of my identity. However, as a vendor I might  
30 like to know that when I execute a query I am not paying to receive duplicated data  
31 records for the same user. One useful way to prevent this is to allow a user to also state



1 with some records “this is the only user profile for which I am choosing to submit this  
2 piece of information”. For example, if user A purchases book X then the user might  
3 choose to associate information about that book purchase with just one of his/her profiles  
4 in the SDI data warehouse, e.g. the profile for the “book-reader”. Stating this allows a  
5 vendor to specify within its query that it is only interested in receiving information  
6 associated with profiles that include this “exclusivity” claim. The SDI client-side proxy  
7 can provide a guarantee that the user only associates data with one profile, for example  
8 using a cryptographic method, e.g. signing the data record with a key to indicate that it is  
9 a unique record.

### 11 **9.3.2 Data Perturbation**

13 In addition to an agent’s identity management and data-release policy an agent must be  
14 careful that it does not compromise its identity-management policy by releasing identifying  
15 information. For example, an agent must be careful not to release the same piece of  
16 information X under pseudonyms P1 and P2 if it is unlikely that the information would  
17 relate to two different agents. In the next section we discuss the concept of data perturbation  
18 in some detail. The idea is that an agent must add enough noise to any information that is  
19 released to prevent identification by another agent, unless the other agent already knows the  
20 agent’s true identity. This date-perturbation module can run on top of identity and data-  
21 release methods. Without careful management of the control of information a  
22 pseudonymous identity management policy is redundant. Remember, whenever  
23 pseudonymity is compromised there is no technical method to prevent vendors and other  
24 agents from exchanging information about the agent.

26 Agents are careful to reveal only information that will not allow a vendor to link the identity  
27 of a user across multiple pseudonyms, defeating the identity management policy.  
28 Pseudonymity can be broken whenever an agent reveals the same piece of information, X, to  
29 multiple agents, e.g. A and B, and that information X has significant discriminative power.  
30 Agents A and B might be able to deduce with quite high probability that it is likely that the



example, my year of birth and Zip code might be almost as valuable for customization purposes as is my full date of birth and street address.

### 9.3.2 Click stream data.

One type of data that is especially interesting in a browsing environment is click stream data, which is stored at the client machine and represents a sequence of clicks that a user has executed, possible across multiple vendors. The data can be stored on a user's local client machine and periodically released under an appropriate pseudonym to the central SDI data warehouse. The client-level proxy server that runs on a user's host machine is in a unique position of being able to monitor the user across different pseudonyms and across different vendors' sites. The client proxy might also collect information about:

- The information that is displayed to a user (e.g. the text, the pictures, etc.)
- Information typed at the keyboard, and profile information transferred from the client machine to a vendor.

The data is gathered by passively observing the actions of the user, and not by direct question-and-response. Possible click stream data release policies include the following data-release policies:

- A. Release no information.
- B. Only release data on the URLs of the most recent sites visited.
- C. Release data about the URLs of the most recent sites visited, and the information displayed to the user.
- D. Release data about the URLs, the information displayed, and the information entered by the user.

in combination with a suitable identity-management policy. For example, an agent might state that all click stream data should be released under the pseudonym that the agent adopts for a particular transaction, i.e. all click stream data with amazon.com should only

1 be associated with my “buying books” pseudonym. An alternative policy might state that  
2 click stream data should be stored under a unique pseudonym, and not identified with any  
3 of a user’s other pseudonyms. The agent might then allow a vendor to perform limited  
4 queries on that data, for example only query information related to particular domain  
5 names. Many variations are possible: e.g. use a unique pseudonym for each new URL  
6 domain, i.e. whenever a user skips to a new site, submit click stream data under a new  
7 pseudonym.

8  
9 Click stream data can be subject to random perturbation, just like standard data-- for  
10 example removing time-stamp information and adding noise to the URLs that a user clicks.

11  
12 There may be other click stream data release policies worth considering outside the  
13 preferred embodiment discussed above. For example, Intermind’s patent number XXXX,  
14 entitled XXXX, provides for the release of a user’s information according to the terms  
15 and conditions of the user’s own data disclosure policy.

16  
17 Such an arrangement could be further enhanced by conditioning the quality of the data  
18 released by the vendor to the user on the strictness of the user’s own data disclosure  
19 policy -- users could then be rewarded to share more of their clickthrough data. Various  
20 levels of security could also be guaranteed to the user releasing such information – for  
21 example the data perturbation technique could be used to ensure that the vendor only  
22 receives aggregate information about his visitors.

#### 23 24 9.4 Agent-Agent Data Management

25  
26 It is important that an agent that submits data to the central SDI data warehouse also  
27 maintains careful control over the data that is directly released to other agents through  
28 one-to-one interactions. It is necessary to prevent a “black-market” in data. Furthermore,  
29 a user might simply prefer that another agent (for example representing a vendor) does  
30 not know certain pieces of information.

1 With respect to preventing a black market in information about a user, it is important to  
 2 prevent two different agents linking the identity of an agent under two different  
 3 pseudonyms, based on comparing information provided by an agent under each  
 4 pseudonym. We do not propose a technical solution to prevent these agents exchanging  
 5 information (outside of SDI) about the agent when this type of linking is possible. Such  
 6 violations can not only affect a user's privacy, but might also decrease the value of  
 7 information provided to the SDI data warehouse, since information can be readily  
 8 exchanged between agents outside of the constraints of SDI.

9  
 10 Fundamentally, the policy under which an agent handles data release to another agent is  
 11 no different from the policy with which data is released to the central SDI data exchange.  
 12 This is a special case, in which the agent knows that it is providing explicit permission  
 13 for an agent (specifically the agent with which it interacts) to receive all information that  
 14 it releases. As indicate above, the pseudonym and the data release to another agent can be  
 15 determined by certificates that the agent is able to present, and an agent's local data and  
 16 identity management policies.

17  
 18 For example, suppose that the policy defines that certificate A confers the right to receive  
 19 a persistent pseudonym, and also receive links to the pseudonyms in set P1. Suppose the  
 20 certificate B confers the right to receive a anonymous pseudonym, and receive no links,  
 21 and certificate C confers the right to receive a persistent pseudonym, and links to  
 22 pseudonyms in set P2. Finally, suppose an anonymous profile is the default. Now,  
 23 suppose a vendor presents certificates A and C. This vendor receives a persistent  
 24 pseudonym, and links to pseudonyms in the union of sets P1 and P2; a vendor that  
 25 presents A and B receives a persistent pseudonym and links to pseudonyms in set P1; and  
 26 a vendor that presents no certificates receives an anonymous profile and no links to other  
 27 pseudonyms.

28  
 29 9.5 Light Clients: Web-Centric Data Management  
 30



1 Figure 12 illustrates a sequence of steps that can allow a light client to execute a  
2 customized interaction with another agent. Suppose agent 1 is a light client. Initially  
3 agent 2 contacts agent 1, requesting profile information so that agent 2 can provide a  
4 customized service to agent 1. Agent 1 receives the request, and contacts the SDI profile  
5 management server, which implements the functionality of an agent's data management  
6 and identity management policies, and provides an identifier for the appropriate agent  
7 pseudonym to agent 2. Agent 2 can then request information about agent 1 directly from  
8 the central SDI server, and finally provide personalized information and services to agent  
9 1.

10  
11 The pending patent application "A System for Location Enhanced ..." [FILL IN NAME  
12 AND PATENT #] describes a method to identify a user based on location, which can  
13 allow automatic detection of close SDI-enabled agents based on physical proximity. The  
14 methods taught in the LEIA patent also suggest the use of a common user identifier,  
15 which could via simple look-up act as the Unique User identifier in the central SDI  
16 database. Other biometric variations, e.g. via Iris scan technology promise to allow  
17 personalized interactions between people and other agents that they approach in the  
18 physical world, with sessions automatically configured based on identification and  
19 principles of consent. Location based filtering can be further used to allow relevant  
20 information to be delivered to a user based upon his/her physical location. It is also  
21 possible to pre-cache personalized information directly to the user's client (and server)  
22 thus overcoming the bandwidth bottleneck or wireless connections.

23  
24 Web-centric profile management can be used in an environment of mobile users that  
25 interact with many different computational devices, for example ATM machines, point-  
26 of-sale terminals, etc. A simple variation of SDI allows a user to allow a machine with  
27 which he/she interacts to be automatically configured for a personal session, based on  
28 carefully controlled queries from that machine to the central SDI data warehouse. A  
29 common "user identifier" be it cards, codes or biometrics can be used to identify and  
30 personalize the local SDI client machine. The user may also be identified during off-line

1 transaction via his/her credit or debit cards, for example on ATM machines and POS  
2 kiosks, providing promotional offers and coupons.

### 3 4 9.6 Smart Cookies

5  
6 Current practice in web-browser based consumer-to-business electronic commerce is to  
7 use cookies, which are identifiers placed on a user's hard drive, to identify a user across  
8 an extended period of time. For example, if I access the New York Times from my home  
9 personal computer on Monday, and then again on Tuesday, the New York Times server  
10 can identify that I am the same individual and build a profile of my interests, i.e. the new  
11 stories which I choose to receive first each day. This type of information about all of the  
12 users that read the Times can allow collaborative-filtering type techniques and  
13 personalization of information in the future, such that my "front page" is different from  
14 the front page of someone else.

15  
16 However, cookies have the unfortunate side effect of allowing an individual to be tracked  
17 across the web pages of different vendors, for example across the web page of the New  
18 York Times ([www.newyorktimes.com](http://www.newyorktimes.com)) and Amazon ([www.amazon.com](http://www.amazon.com)), if the Times  
19 and Amazon both embed content from the same third party in their pages. This happens,  
20 for example, with the DoubleClick advertising network. DoubleClick  
21 ([www.doubleclick.com](http://www.doubleclick.com)) operate a virtual network of pages, and can track a user across  
22 any page within their network, and gather a very comprehensive user profile. Although  
23 the cookie mechanism is designed so that only vendors with the same domain name can  
24 access cookies on a user's hard drive, they can easily be used to profile users across  
25 multiple vendors, for example with a double-click style network that embeds a universal  
26 advert server within each page.

27  
28 The system of SDI allows controlled personalization, such that a proxy-automated log-in  
29 session where the proxy presents a user's pseudonym to a vendor allows that vendor to  
30 track a user over time at its own web site, but a vendor cannot track a user across web  
31 sites. The SDI client-side proxy agent will disable cookies in their current form (although





1 BGGMM98;GGMM98] or W3C proposals [CR98; RC99; W3C-OPS 97]; essentially the  
2 pseudonym and an associated password that a user adopts for a vendor allows log-in to  
3 that vendor.

4  
5 Also important when supporting a system that allows a user to browse pseudonymously  
6 is that the physical attributes of a network system are removed of their identifying  
7 characteristics. For example, another role of client-side data management agents is to  
8 strip the 'from' field in a HTTP/TCP message.

## 9 10 9.7 Implementation Details

11  
12 In an Internet browser environment the client-side proxy agent that provides distributed data  
13 management for an agent might be implemented as a plug-in into the browser, that can for  
14 example be downloaded from a central SDI server. The browser is then configured to use  
15 the SDI proxy as its proxy, and the SDI proxy itself connects through a user's ISP (or other  
16 intranet gateway) to the Internet, and on to other vendors.

17  
18 The user provides his/her SDI proxy with personal information, such as his/her name,  
19 mailing address, and e-mail address. The client-level proxy registers then registers the user  
20 with the central SDI server, providing the server with the name, address and e-mail address  
21 of the user. Other basic user information might include demographic information, for  
22 example a users job, marital status etc. The client proceeds to automatically generate a  
23 unique SDI user ID code, and a private key to allow future authentication of its log-in.

24  
25 The client can create a unique public key/private key pair. This key pair can be generated  
26 only once for a person, and although the central SDI user ID server does not know the key  
27 pair, the server can verify that a key pair is only generated once-- because a new user must  
28 present proof of identity to establish an account. The client generates a unique user  
29 identifier, UUID, for example with the methods taught in [Chaum 85; Schneier 92]. The  
30 UUID can then be blinded and signed to certify that a user is registered with SDI, using  
31 Chaum's technique of blinded signatures so that the certifying agent does not the identity of

The client-level proxy can now sign messages with its private-key, and provide the signed to UUID, to verify that (1) the UUID represents a validated user; (2) it is the client-level proxy authorized to act for the user, because it has the private-key associated with the UUID. The client-level SDI proxy uses the private key to authenticate messages that it sends to other modules within SDI, such as Pseudonym administering servers. The unique user ID for a user does not carry any information about the user, its sole purpose is to provide a unique identity. The unique UUID can also be used to generate new pseudonyms for users that are certified one-time for a particular vendor, so that the user certifies that he/she maintains the same profile for all interactions with a vendor.

At this stage the central SDI server might verify the identity of the user, and also check that the user is not already registered with SDI. The method for verifying the identity of a user could include requesting that the user provides his/her social security number, or some other institutional solution that is used for this purpose. In the future we could envisage an electronic system for such an identity procedure, but the method might require for the user to execute this initial step in person with the presentation of a recognized photo ID. The central SDI user ID server maintains a database of all users that are registered with SDI, and checks that that the user is not already registered with the system of secure data interchange. The central SDI proxy may also provides the user with a signed certificate of some universal identifier, such as its Social Security Number, that the user can use to generate other certificates from certifying agencies, to be able to gain certificates under pseudonymous identities.

### 9.7.1 Initializing Identity and Data Management Policies

The client-side SDI proxy might provide a rule-based interface to allow a user to select from a menu of defaults an appropriate profile management policy, or to configure with

09699098-102700

1 the aid of a decision tree an appropriate profile management policy for a user. Profile  
2 management policies define how a user will interact with various classes of vendors  
3 (depending on the nature of the business that the vendor is engaged in), the kinds of uses  
4 to which the transactional information that a vendor collects can be put to, and the  
5 amount of information that a vendor which collects profile information about a user is  
6 authorized to release. The client-level proxy manages a user's interactions with vendors,  
7 to keep them within desired policies.

8  
9 **9.7.2 Generating a New Pseudonym**

10  
11 The method of blinded signatures [Chaum 85; Chaum 92] provides a useful technique to  
12 generate new pseudonyms for a user, without any centralized database that stores  
13 pseudonyms. The client-level SDI proxy can generate a new identifier, comprised of a  
14 sequence of bits, that will be unique with a high degree of probability. The identifier is  
15 "blinded" and then submitted to a trusted-third party to be authenticated for use as a  
16 pseudonymous identifier. When authenticated the blinding factor is removed, and the final  
17 signed pseudonymous ID can be used as a new identifier. The signature can allow a user to  
18 associate certificates with that identifier, and also to verify that the user has only a single  
19 pseudonymous ID with a particular vendor.

20  
21 Pseudonym administering authorities (PAS) cannot build dossiers of the pseudonyms,  
22 because users submit "blinded" identifiers. The only information that a PAS has is the list of  
23 unique vendors that a particular user has registered with. The pseudonym administering  
24 server can be operated by an agent with a trusted relationship with a particular vendor.

25  
26 Every SDI user has a unique identifier, a UUID, that is presented to the PAS with the new  
27 pseudonymous identifier to be validated. The PAS can verify that this is the first pseudonym  
28 for a particular vendor. Each Pseudonym administering server has a public key / private key  
29 pair (PKPAS, SKPAS) for each Vendor for which it validates new pseudonyms. A signed  
30 pseudonymous identifier, signed with the private key of PAS, verifies that the pseudonym is  
31 valid. Cryptographic techniques ensure that the signature cannot be falsified, and allow



The proxy agent then implements the profile-management policy, as it relates to the certificates presented by the vendor. First, the agent might submit the basic profile information to the vendor, for example a user's age, nationality, state, sex—anything a user is happy for any vendor to know. This is the basic profile that is configured by the user during initial registration with SDI. Later, the vendor might request profile information, that can be provided if the vendor is authorized to receive such information.

### 9.7.4 Leveraging Existing Standards

The architectural framework outline above can be implemented with a number of existing technical methods. One approach is to use the Extensible Markup Language (W3C-XML) to encode information exchanged between client agents and vendor agents, for example profile information and requests for information. This has been suggested by the W3C consortium. The eXtensible Markup Language (XML) proposal of the Worldwide Web Consortium working group on SGML provides an ideal standard for representing such information. XML allows meta-content to be included with documents, machine-readable information that enables documents to be processed by client software. Augmenting web documents with structured information enables clients to respond to requests for profile information. XML can represent rich data structures, and that allows a grammar to be defined for information that allows data to be automatically verified for correctness.

### Example: An XML Representation for a Profile Request Message

2 A vendor might request profile information from a client-side SDI proxy agent. A vendor  
3 sends a 'Request Profile' message, and the client-side SDI proxy agent responds with a  
4 'Profile Update' message, that contains profile information, in accordance with a user's  
5 profile management policy. The request-response mechanism can be implemented using the  
6 standard HTTP Post/Response mechanism in conjunction with XML message types. The  
7 'Request Profile' message can be represented in XML as:

```

9 <?XML version = "1.0"?>
10 <?xml:namespace ns = "http://www.sdi.com" prefix = "SDI"
11 ?>
12 <!doc>
13 <SDI:Request> http://www.some_vendor.com </SDI:Request>

```

15 and the SDI proxy agent's 'Profile Response' message can be represented in XML as:

```
17 <?XML version = "1.0"?>
18 <?xml:namespace ns = "http://www.sdi.com/clientX12345"
19 prefix = "SDI" ?>
20 <!doc>
21 <SDI:Update>
22 <SDI:Profile>
23 <SDI:Item> (1231, 0.453) </SDI:Item>
24 <SDI:Item> (1041, 0.034) </SDI:Item>
25 </SDI:Profile>
26 </SDI:Update>
```

28 An illustrative Document Data Type (DTD) for an SDI:Profile element type is presented in  
29 the next section. The XML messages are included in the body of standard HTTP  
30 Post/Response messages. We limit the performance degradation caused by out-of-date  
31 profile information that is stored within web pages of on-line vendors by associating “out-

The system as outlined above can be implemented within the current HyperText Transfer Protocol (HTTP), as a sequence of challenge/response pairs between clients and servers. The HTTP Post/Response mechanism allows clients and servers to exchange data, and this data can be an instance of an XML Document Type, within the body of a HTTP message. The HTTP protocol is the underlying mechanism, with SDI messages contained in the body of the HTTP Post and HTTP Response as XML documents.

### Example: A Possible XML Representation of a User profile

The World Wide Web Consortium (W3C) SGML working group developed XML (extensible markup language) to provide an open and extensible grammar for structured data [XML]. An XML document has an associated schema definition to enable an XML-enabled browser to validate the structure of XML data automatically. A Schema in XML is called a Document Type Definition (DTD), and defines the names of tags, their structure, and their content model. XML allows the DTD for an XML file to be identified through a Universal Resource Indicator [URI] in the header of the file (see below). XML also allows URIs for mobile code resources to be referenced, in order to enable a client to process embedded XML data. An XML document must be well formed, and in order to be well formed the tags must form a tree structure. In addition, the DTD allows the structure of an XML document (an instance) to be validated against a particular schema. Senders and receivers must only send valid SDI files. Each SDI message is a valid XML document.

We provide an example XML instance and part of a Document Type Definition for use within the system of SDI. We assume in this example that profile information is represented





The Document Type Definitions for this document are specified in the header, and include URIs to a DTD of the Open Profiling Proposal of the W3C, and also a DTD of the Secure Data Interchange. The OPS DTD is used to boot strap the SDI DTD, providing tags for common profile information, such as 'Gender', 'Age', 'Income', etc. The section of the SDI Document Type Definition that is used in the above XML fragment is presented below. It makes reference to tags defined in the OPS DTD, and the RDF (Resource Description Framework), a W3C proposal to standardize the structure of Digital Type Definitions for XML documents. XML Name spaces [NS] provide a method for unambiguously identifying the semantics and conventions governing the particular use of property-types by uniquely identifying the governing authority of the vocabulary, for example OPS and SDI in the example above. The URI for a schema can contain a human and machine-readable description of an XML schema.

```
<!ELEMENT SDI:ProfileData (SDI:Location?,
OPS:Demographic?, SDI:ID?, SDI:Profile?) >
<!ELEMENT SDI:Location (SDI:Geocode?, SDI:DigiMap,
OPS:Zip?, OPS:Address?) >
<!ELEMENT SDI:ID (OPS:Name?, SDI:PublicKey?,
SDI:Pseudonym?) >
<!ELEMENT Profile RDF:list<SDI:Profile-item> >
<!ELEMENT SDI:Geocode #PCDATA >
<!ELEMENT SDI:Digimap #URI >
<!ELEMENT SDI:PubicKey #PCDATA >
<!ELEMENT SDI:Pseudonym #PCDATA >
<!ELEMENT SDI:Profile-item (SDI:Attribute-ID,
SDI:Attribute-value) >
<!ELEMENT SDI:Attribute-ID #PCDATA >
<!ELEMENT SDI:Attribte-value #PCDATA >
```

The tag '#PCDATA' is used here to represent numeric or textual information, '#URI' declares that an instance of element 'SDI:Digimap' must be a valid URI pointer.







A. **[Strong Protection]** Assume that the adversary knows that an agent **A\_1** with data records **P** submits data into a database; i.e. assume that the adversary knows for sure that one of the data records in the database relates to agent **A\_1**.

Strong protection is only possible if there are enough data records from other agents to allow agent **A\_1** to add noise to its own data and make it sufficiently like that one of the other data records is the data of **A\_1**; in particular the other data records must also be submitted with random noise perturbation, or rounded, such that it is possible that other data records are submitted by agent **A\_1**.

B. **[Weak Protection]** Assume that the adversary does not know that the agent A<sub>1</sub> with data records **P** has definitely submitted data into the database.

Weak protection is easier to achieve. It is only necessary to add enough random perturbation to data to make the number of possible data records over the population of possible data records that are supported with the perturbed data record large enough to prevent agent identification.

Randomized data is still useful within SDI for data mining and other applications, so long as the amount of noise which is added to records is small in comparison to the value of a record. For example, we can still perform correlation across fields with randomization so long as the randomization does not destroy any trends between fields. Randomized data is marked as such within SDI, and labeled with the degree of degradation, so that SDI can be aware of the number of records to get relevant accuracy levels, and can report accuracy to customers. There is a tradeoff between the level of privacy protection and the level of aggregation at which responses to queries become accurate. It is possible to add random noise to data but still allow data that is aggregated across multiple records to be quite accurate, so that useful data mining can be performed. Binning or rounding of data does not have the same effect. For example, assuming additive noise and additive aggregation, then randomized data from a number of agents can be aggregated to obtain an aggregate value







For example, suppose the distribution is  $\mathbf{f} = \mathbf{N}(\mathbf{0}, \mathbf{1})$ , i.e. a Normal distribution with mean 0 and standard deviation 1 and the value  $\mathbf{X} = \mathbf{1}$ . Assume that the random number generated by  $\mathbf{f}$ ,  $\mathbf{noise} = +\mathbf{0.3}$ . Therefore the perturbed value is  $\mathbf{X}' = \mathbf{1.3}$ , and this is the information that the agent submits to the database. Given this, an adversary can compute that a ex post distribution for the true value,  $\mathbf{g}(\mathbf{X}) = \mathbf{N}(\mathbf{1.3}, \mathbf{1})$ .

Bayesian analysis can be used to compute an optimal parameterization for the random distribution, given a desired level of protection and information about the data already in the database, and the randomized perturbations that were added to that data. Assume that the database contains a set of data records submitted by agents, each with the submitted value  $\mathbf{X}'$  and the noise distribution from which it was generated,  $\mathbf{f}$ . With this information agent **A 1** can select an amount of noise perturbation.

The goal is to compute  $\Pr(\text{Agent} = \text{A\_1} \mid \text{true value} = \text{X})$ , i.e. the probability that the agent that submits (**perturbed value** =  $\text{X}'$ , **noise distr.** =  $\text{f}$ ), is the agent with true value  $\text{X}$ . Using Bayes rule, this conditional probability is computed as  $\Pr(\text{A\_1}) \times \Pr(\text{true value} = \text{X} \mid \text{A\_1}) / \Pr(\text{true-value} = \text{X})$ . Each of these terms can be computed as follows:

**Pr(A\_1)** This is the a priori probability that a random data record is submitted by agent **A\_1**, and is equal to  $1/N$ , for  $N$  records.

**$\Pr(\text{true-value} = \mathbf{X} \mid \mathbf{A\_1})$**  This is the probability that an agent that submits  $(\mathbf{X}', \mathbf{f})$  has true value  $\mathbf{X}$ , and is computed from  $\mathbf{X}' = \mathbf{X} + \mathbf{f}(\mathbf{d})$  as  $\mathbf{g}(\mathbf{X}) = \mathbf{X}' - \mathbf{f}$ .

**Pr(true-value = X)** This is computed as the sum of  $\text{Pr}(\mathbf{A}_i) \times \text{Pr}(\mathbf{X} \mid \mathbf{A}_i)$  over all data records, and represents the probability that any one of the data records represents data with true value **X**.

Agent **A\_1** can now select parameters for distribution **f** based on this analysis, so that the probability that it is the agent to submit the new data record is less than **eps**, its desired protection level.

The rule has the right behavior—the more data present in the system then the more accuracy an agent can use to submit its own information, for the same privacy level  $\epsilon$ . Notice that the agent is more protected as:

1.  $\Pr(\mathbf{A\_1})$  decreases, i.e. with more data records in the database
2.  $\Pr(\text{true-value} = \mathbf{X} / \mathbf{X'}, \mathbf{f})$  decreases, i.e. with more noise perturbation  $\mathbf{f}$
3.  $\Pr(\text{true value} = \mathbf{X})$  increases, i.e. there are more data points from other agents that might have true value  $\mathbf{X}$ .

The parameters for  $\mathbf{f}$ , for example the mean and standard deviation in the case of noise generated from a Normal distribution, can be selected to set  $\Pr(\mathbf{X}' \mid \text{true-value} = \mathbf{X}, \mathbf{f}) = \mathbf{eps}$  for the case that  $\mathbf{X}' = \mathbf{X}$ , i.e. when the random noise is zero. This presents the worst-case, assuming that distribution  $\mathbf{f}$  places the most probability on zero noise.

Finally, to compute parameters for  $\mathbf{f}$  the agent requires aggregate information about the data in the database that has been submitted by other agents. In particular, the agent needs  $\Pr(\mathbf{X})$ , the probability that any of the current data records could correspond to data with true value  $\mathbf{X}$ . In some cases it is important to receive this information without revealing true value  $\mathbf{X}$ .

In a simple case, for example when a trusted intermediary such as SDI stores information provided by agents, the agent can simply poll SDI for  $\Pr(\mathbf{X})$  before selecting a level of noise to use to perturb its information. Alternatively, when SDI is computing the level of noise perturbation to add dynamically as information is released in response to queries then the

1 system of SDI can compute  $\Pr(X)$  directly. In a more complex case, consider a problem  
2 where an agent is releasing information directly to an adversary, and the relevant set of data  
3 points are data that the adversary agent has already collected from other agents. In this case  
4 the provision of information  $\Pr(X)$  must be done within a secret-protocol where the  
5 adversary does not learn the value of  $X$  in the process. A straightforward way to achieve this  
6 is for the “adversary” agent to post aggregated information about the probability that a data  
7 record in its population has true value  $X$ , for a range of different values of  $X$ , and allow an  
8 agent to use anonymous look-up in a table.

9  
10 In summary, the following procedure can be used to select an appropriate level of noise to  
11 add to a data point:

- 12
- 13 a. Choose a distribution family (e.g. Normal, Uniform, etc.), and let  $\text{Par}$  denote the  
14 parameters that define a specific distribution.
- 15 b. Choose a level of privacy protection,  $\epsilon$  where  $0 < \epsilon < 1$ .
- 16 c. Request  $\Pr(X)$  from the database, i.e. the current probability that a perturbed (or  
17 otherwise) data record in the database has true value  $X$ .
- 18 d. Compute parameters for the distribution to set  $\epsilon = \max_{\{X'\}} \Pr(A_1 | X', X, f)$   
19 where  $X'$  is the value generated from  $f$ .
- 20

### 21 10.1.2 Examples: Strong Protection

#### 22 (a) Uniform Additive Noise Perturbation Distributions

23 Assume that every agent submits information from a uniform noise distribution, centered  
24 around its true value. The decision variable in choosing a level of data perturbation when  
25 submitting new information is the range of the uniform distribution.

26  
27  
28  
29 Suppose agent 1 submits data point  $g_1 = [4, 6]$ , to denote that its perturbed value is  $X'$ , and  
30 the value was computed with additive uniform noise  $U(-1, 1)$ . Agent 2 submits data point

1  $g_2 = [4,6]$ , agent 3 submits  $g_3 = [3,5]$ , agent 4 submits  $g_4 = [3,6]$ , and agent 5 submits  
2  $g_5 = [2,4]$ .

3

4 Suppose that an adversary knows that user John has true value  $X = 2.5$ , and that an agent for  
5 John has submitted a data record.

6

7 In this case the adversary can be sure that agent 5 represents John, because  $\Pr(A_5 | x = 2.5)$   
8  $= \Pr(A_5) \Pr(x = 2.5 | [2,4]) / \Pr(x = 2.5) = 0.2 * (1/2) / (0.2 * 1/2) = 1$ .

9

10 In comparison, for an adversary that knows that Mary has true value  $X = 5$ , then  $\Pr(A_1 | x$   
11  $= 5) = 0.2 * (1/2) / (0.2 * 1/2 + 0.2 * 1/2 + 0.2 * 1/3) = 0.375$ , and similarly for agent  $A_2$ . The  
12 probability that agent  $A_4$  is Mary is 0.25.

13

14 Now, a new user, Bill, with agent  $A_6$  wants to provide information about its data  $X = 5$ ,  
15 and wants to be sure that an adversary cannot determine its identity with probability greater  
16 than  $\epsilon = 0.1$ . Let  $d_6$  denote the uncertainty selected by agent  $A_6$ , and compute an  
17 optimal  $d_6$  as so that  $\Pr(A_6 | x = 5)$ , i.e. the probability that agent  $A_6$  has true value  $x =$   
18 5. This is computed as  $\Pr(A_6 | x = 5) = 1/6 * (1/d_6) / (1/6 * 1/2 + 1/6 * 1/2 + 1/6 * 1/3 +$   
19  $1/6 * 1/d_6) = 1 / (4/3 * d_6 + 1)$ , and is less than  $\epsilon$  for  $d_6 > 6.75$ . Therefore agent  $A_6$   
20 should generate a perturbed value with additive noise computed with uncertainty 6.75, and  
21 can then be sure that its identity is protected.

22

23 (b) Normal Additive Noise Perturbation Distributions

24

25 Now, assume a normal noise distribution, and let  $sd_i$  denote the standard deviation selected  
26 by agent  $A_i$  for its noise distribution. The analysis is slightly more complicated, because  
27 the  $\Pr(X | A_i)$  is now a function of the position of  $X$  within the distribution, not just  
28 whether it is in range as with uniform distributions.

29

30 In this case an agent chooses the standard deviation for its Normal distribution by assuming  
31 that the randomized value  $X' = X$ , i.e. that the random noise distribution generates zero





1 This technique of binning information, or rounding information, is simpler to implement  
2 that random noise perturbation because probability distributions must neither be sent to the  
3 provider or maintained by the provider. However, the information is represented with a  
4 static accuracy which is hard to improve as the amount of data submitted increases. This  
5 occurs automatically in the standard version.

### 7 10.3 Adding Random Perturbations to Discrete Data

9 Adding noise to discrete values is a little more complicated, because the new value must  
10 remain feasible. For example, in randomizing the name of a CD the new name must be the  
11 name of another CD, not some "made up" name. We need to add noise to make data  
12 elements "close" to the accurate values. With discrete data, such as the name of an artist,  
13 "close" must be defined within the correct metric. The appropriate metric is such that a  
14 "close" value shares many of the same characteristics. For example, it is not appropriate to  
15 assign a close value on the basis of a shared last letter in the first name, but it is appropriate  
16 to assign a close value on the basis of an artist from the same genre of music --- from a  
17 "semantic cluster".

18 Agents use discrete probability distributions to randomize data points.

19  
20 For example, suppose that an agent wishes to reveal the name of the artist that recorded the  
21 last compact disk that a consumer purchased. An artist's name is best viewed as a point in  
22 "artist space", and therefore as a discrete value. It does not make sense to change a random  
23 set of letters, because the new "name" will not be the name of a valid artist. Instead, the  
24 concept of noise is to randomly choose a new name close to the current name, i.e. choose a  
25 new feasible location in artist space. One reasonable solution for the names of recording  
26 artists is to define a neighborhood of artists that are close to the original artist, and select a  
27 new artist from within the neighborhood with equal probability. The metric that defines how  
28 close artists must be computed using a system that is common knowledge to the provider of  
29 information and the consumer agent, because the provider of information must be able to  
30 compute the distribution over true artist name, given a randomized name and probability  
31 distribution. We can define the probability distribution with a threshold in the distance









We suggest a simple technique to bootstrap the system. The method works in the case that there are a number of different reasons to decline to provide information to another agent in response to its query. An agent provides randomized data and a privacy level  $\epsilon$  to the database, e.g.  $([16, 20], 0.2)$ , where  $[16, 20]$  is the range of possibilities for an agent’s true age, and  $0.2$  is the desired privacy level.

11 The SDI database can be configured to only release information when it can verify that it is  
12 safe to. Similarly, an automatic method on the client machine of a user can be configured to  
13 only release information when it is safe to do so. The problem that is solved here is that  
14 sometimes there is no reasonable amount of noise that can be added to a data value to  
15 protect a user's identity.

Another technique to solve the bootstrapping problem could hold all data until it is safe to release information for the majority of data points, given the state privacy requirements of users etc. The measure of “safe” is defined by the information in the database and the privacy level specified by each submitting agent. A more advanced technique could introduce data periodically as new parts of the data space become populated, so that there is never an “active” under-populated part of the space.

## 24 10.6 Verifying Noise Levels with Playback

26 We can use a cryptographic technique to verify the distribution of noise that is added to data  
27 - and also to enable replay. “Playback ability” - the ability to reconstruct the original record  
28 from a noisy version of that record is important for a number of purposes. An individual  
29 may want to obtain proof of a transaction for legal purposes and law enforcement agencies  
30 with appropriate warrants might want to examine original records.



integrated into an existing advertising service which they operate and maintain, in which case the vendor receives an appropriate fee for data which is exchanged between his/her existing customers, and a reduced fee (which may be split with the central SDI service) for data which is exchanged by/between a member of his/her SDI service and vendors who are members of the central SDI service but not of his/her local SDI service. Unless or until the client-level proxy server becomes a standard or a large critical mass of end-users adopts the service, one of several compelling business models could be used by a vendor to encourage the visitors to his/her site to adopt the user-centric SDI service based upon the monetary incentives the user may receive for subscribing. Typically the resulting revenues are split between the user, the user-centric SDI service and the vendor (in exchange for promoting the service to their site visitors a share of resulting revenues generated may be necessary).

#### 11.1 An open SDI system

An ISP level proxy server can contain the user profile generation module, profile processing module, user profile interest summary generation module and target object generation module which operate in distributed manner. This enables an ISP to independently implement the core functionality of the system without the cooperation of information vendors (Web sites) or their operators (Web hosts) who opt-out of SDI. The modules in third-party SDI servers can share information with the modules in network vendor servers. This flexible architecture enables the user-centric SDI service to be implemented (by ISPs or completed independently) and when available inter-operating with the complete data sets available from the information vendors.

SDI can allow third parties to operate their own secure advertising and/or electronic commerce-based product syndication affiliate network (for all customers). In accordance with the preferred implementation of SDI, these "advertisements" are represented at the item-level and may be distributed across the (proprietary or main) SDI system network presented in the form of ads, affiliate or portal links to purchasables or sites (which may contain target objects as purchasables) and/or simply (transparently integrated) selections

on an e-commerce product catalogue. The users that are also subscribed to SDI, can be given highly personalized information for each site or for the network of sites (which could involve an interface which provides site to site links as a “virtual mall”), and a menu interface to these sites which includes the 2 or 3 dimensional personalized menu features and personalized search facilities as disclosed in the parent description (a “personalized portal”). Traditionally the term “portal” was exclusively reserved for major search engine/directories such as Yahoo, Lycos, Alta Vista etc.

9 However, at the present time, there is a trend in which many Web-sites are becoming what  
10 are called “portals” or “community portals” serving a particular e-commerce or content  
11 niche or theme. Often highly robust content and/or extensive (usually web-wide) links to  
12 relevant information resources and web-sites are accessible from such portals in order to  
13 create a “one-stop shop” service to visitors or community members. Additionally, in  
14 order to leverage low-cost access to more robust content, a transaction-based or “click-  
15 through-based” or transaction based model is also emerging. This model fits particularly  
16 well within the current framework of SDI. In particular user profiles provided by SDI-  
17 enabled site visitors provides the platform for automatic generation of not only  
18 personalized ad banners but also content and portal links, it is reasonable for these  
19 community portals as part of their business relationships with the sites and advertisers to  
20 which they provide access to share revenues paid by these sites and advertisers with the  
21 end-users (which significantly increase the click through by virtue of personalization of  
22 information at potentially all varieties at the site resulting from revelation of the profiles  
23 from those visitors). As discussed on the parent case, these pages are pre-cached in  
24 advance onto the local server and client.

26 Within today's internet infrastructure as pages are typically dynamically generated, the  
27 files are first pre-cached then dynamically generated local to where they are anticipated  
28 to be used. Additionally, high-end community portals typically offer content from  
29 external sources which is sometimes syndicated from content providers which are paid by  
30 click through or transaction. In this model the user or the site may absorb these fees paid  
31 to the content provider (because the content when personalized may appeal to almost any







## 19 12. Ancillary Systems

21 12.1 Support for Pseudonymous electronic mail

23 The ISP-level proxy server is positioned just behind the firewall of the user's local dial-up  
24 network (ISP or Intranet). The proxy provides protection for users operating under  
25 pseudonyms from point-to-point attacks and HTTP header-tracking by stripping HTTP  
26 header-information and forwarding HTTP packets on to their destination with no  
27 information other than their source at the ISP-level proxy server. The ISP-level proxy also  
28 supports pseudonymous e-mail, between users, and between users and vendors.

1 Figure 2 shows a couple of users connected to clients, that are in turn connected to the  
2 Internet through a local intranet, such as the network of an Internet Service Provider (ISP).

3 The proxy "washes" outgoing messages of any information that would compromise a user's  
4 pseudonymity, for example the "referral" field that contains the previous URL of a user in a  
5 HTTP message. HTTP messages also leak other information, for example browser software  
6 on a user's client machine, the operating system and a user's IP address.

7 A user can receive electronic mail through the PID and associated IP address of the ISP-  
8 level proxy server.

9 The preferred implementation of this system allows the user to periodically check for new  
10 mail. The client-level proxy gains access to the mail box that is associated with a  
11 pseudonym by providing a correct response (signature) to an ISP-generated challenge.  
12 Notice that with this solution, the ISP-level proxy has no way to connect the pseudonyms of  
13 a user, so long as the user's client is not identified in its messages to the ISP-level proxy  
14 server other than by the PID that the proxy makes a request for.

15 We can extend this mechanism using a technique taught in the Lucent Personalized Web  
16 Assistant (LPWA). The LPWA [BGGMM 97; BGGMM98] provides for a sequential access  
17 mechanism to the mailboxes that belong to a user through a one-way function that takes the  
18 user's SDI log-in name and password, and an integer from 1 to N, and computes the  
19 mailbox location. The mail server does not need to maintain a list of pseudonyms for each  
20 user, because the user is able to efficiently access all of its mailboxes sequentially as a  
21 function of other information.

1 Another variation, that relies on the user placing trust in the ISP-level proxy server, provides  
2 the ISP-level proxy with the e-mail address for each pseudonym. This push method is more  
3 efficient, because the ISP proxy and the client proxy communicate only when new  
4 messages arrive, but provides the ISP proxy with information to compute all the  
5 pseudonyms for a single user—probably undesirable.

6

## 7 12.2 Support for Pseudonymous Physical Mail

8

### 9 12.2.1 Vendor to User

10

11 A vendor must hold a “physical mail certificate” to be able to send mail (packages,  
12 letters) to a user under a pseudonym. The certificate is similar to the “electronic mail  
13 certificate”, in that it is signed by the private key of the user’s pseudonym, and indicates  
14 that the vendor with public key  $P^*V$  can send mail to the user (under the pseudonym).

15 Each user has a trusted physical address authority, just as it has a trusted electronic mail  
16 authority (the second-level proxy server), that maintains the physical mailing address for  
17 each pseudonym. When a vendor has a letter  $X$  to mail to user with public key  $PKP$ , the  
18 vendor generates a unique ID for the package,  $IDX$ , and sends the ID code and the  
19 physical mail certificate to the trusted physical address authority of the user.

20 The physical address authority receives the certificate,  $S(PKP, PK^*V, SEND\_MAIL)$ ,  
21  $SKP$ ), that indicates that the vendor is authorized to send mail to the pseudonym, and the  
22 packages identify code, signed by the vendor to certify that the vendor holds the secret  
23 key that matches the public key in the physical mail certificate.

24 The vendor then passes the letter  $X$  and the signed ID code to a trusted mailer, that  
25 supports pseudonymous mailing, and has been certified by the central SDI server as such.  
26 The trusted mailer then provides the signed ID code to the physical address authority,  
27 signed with the private key of the trusted mailer. The physical address authority verifies  
28 that the trusted mailer is a valid service, and releases the real address of the user to the  
29 mailer. The mailer now has the letter  $X$  that the vendor wants to send to the user with  
30 pseudonym  $P$ , and the physical mailing address of the user - and the package can be  
31 mailed. At not time did the vendor determine the true mailing address of the user, unless  
32 it works in collusion with the trusted mailer, but the trusted mailer is certified by SDI,  
33 and also audited by the chosen physical address authority of the user. The address  
34 authority will only release addresses to reputable pseudonymous physical mail agents.

35 We can operate physical mailing lists in the same way, and gain additional security by  
36 never releasing the pseudonyms or the mailing addresses to the vendor that has requested

1 the targeted solicitations. We can use a technique that is similar to the technique that we  
2 used for virtual mailing lists. The vendor describes its solicitation to the central Secure  
3 Data Interchange, which leverages as much data as possible (without violating the  
4 privacy policies of any of the users or vendors that are represented within the data). The  
5 central SDI server generates a list of suitable pseudonyms, and then provides a series of  
6 unique codes to the vendor, that the vendor can supply to its chosen pseudonymous  
7 mailer with the material that is to be mailed. The central SDI server also provides the  
8 appropriate address authorities with authorization to release the physical mail addresses  
9 to the mailer when presented with the IDs. Notice that at no stage did the vendor have the  
10 pseudonyms or the mailing addresses. The parties all have only as much information as is  
11 necessary - the vendor needs some way to identify its packages to the pseudonymous  
12 mailer. The mailer needs an identifier to present to the address authority, and receives the  
13 addresses. The address authority just needs to know what addresses to release and to  
14 which third party.

## 15 **12.2.2 User to Vendor mail**

16  
17 The Secure Data Interchange system also provides a mechanism for users to send  
18 physical mail to vendors that are registered with SDI with pseudonymous return  
19 addresses. In particular, when a user sends mail to a vendor, the first-level proxy server  
20 provides a tool that: (1) Computes/Looks-up the appropriate pseudonym for the user with  
21 this vendor. (2) Generates a unique ID, and submits a signed message to the central SDI-  
22 server, where the message relates the pseudonym, the vendor, and the ID. (3) Provides  
23 the unique ID to the user.

24 The user writes the unique ID on the envelope, and mails it to the vendor. Should the  
25 vendor wish to reply to the user, then the vendor can take the envelope to a  
26 pseudonymous mailer, and request that the envelope be mailed appropriately. The  
27 pseudonymous mailer verifies the identity of the vendor, and then submits the ID,  
28 together with the vendor's signature, and its own signature, to the physical address  
29 authority that is maintained by SDI. SDI releases the address to the mailer that can then  
30 return the mail.

## 31 **12.3 Pseudonymous Payment Mechanisms**

32  
33 The Secure Data Interchange architecture must be able to support all the standard electronic  
34 commerce functions that we take for granted, but while maintaining pseudonymity for users  
35 and following privacy policies. There are various different solutions to this problem.

### 37 **12.3.1 Anonymous Credit Card Payment [LMP 94]**



## 1 12.4 Client-Side SDI Proxy

2

3 The client-level SDI proxy, implemented as a client program running on the user's client  
4 machine, manages all data transfer between the client machine (and the user), and other  
5 vendors and the central SDI data warehouse. A key function of the client-level proxy is to  
6 implement profile management for a user, to control the ability of agents to track a user as  
7 he/she interacts with multiple vendors. The client-level proxy also controls release of profile  
8 information: the addition of demographic and other personal information to profiles, and the  
9 control of random perturbation to fields to prevent linking across user profiles.

10 The client-level proxy maintains profile information for a user's collection of pseudonyms,  
11 and allows the user to view and challenge profile information. The proxy also provides a  
12 rule-based interface to allow a user to select appropriate privacy/personalization policies.

13 The primary mechanism that protects the identity of a user across multiple vendors and  
14 service providers is the ability to interact pseudonymously with vendors. The user can  
15 choose a unique pseudonym for each third party with which he/she interacts, and be  
16 absolutely certain that he/she is the only party that knows his/her true identity. There is no  
17 way that a vendor can know anything about the transactions that a user has had with other  
18 vendors under alternate pseudonyms unless the user chooses to disclose the equivalence of  
19 pseudonyms, or use the same pseudonym across multiple vendors.

20 It is useful to distinguish three key modes of use of the Secure Data Interchange system:

21

- 22 • Static data-mining. Query execution with no dynamic requests for new  
23 information from information providing agents.
- 24 • Interactive data-mining. Query execution that includes dynamic attempts to  
25 request additional information from information providing agents.

- Client-side data-mining. Query execution that is performed on the local client machine of an agent, based on data stored exclusively on that machine.

All three modes may use distributed information, i.e. it is possible that the information is stored in the central SDI database, or on distributed client-side information servers, or in third-party servers. The first 'static' case and the second 'interactive' case are distinguished from the third 'client-side' case in that the query execution is performed centrally in the SDI data warehouse in the former, and on an agent's client machine itself in the latter. Client-side data mining has particular application to privacy-protected customization of information and services in on-line business-to-consumer applications.

## 1. Static Data Mining

In static data mining queries are executed on the information that is currently present in the SDI data warehouse, and there is no opportunity to contact agents and request more information. Applications of this type of static data mining include all types of "standard" database queries, where it is assumed that the data set is static. Queries may be open-ended, i.e. "find me all data records of this type, and perform the following operations.." or closed, i.e. "perform the following query on the data record for agent with pseudonym P1". Queries may also have side-effects, i.e. "find all records that satisfy this constraint, and then take action A".

## 2. Interactive Data Mining

In interactive data mining multiple agents may be contacted by the system of Secure Data Interchange in the process of executing a query, to request new information from agents, or push information to agents. A central application of this mode of data mining is *matchmaking*, which is a process where information flows between agents if and only if both the profiles of both agents are mutually compatible. In interactive data mining the querying agent does not need to know the identities of agents that are contacted by SDI, this is all transparent, and hidden from the querying agent.

### 3. Client-Side Data Mining

In client-side data mining the querying agent executes a query with the information associated with a specific agent, and the information remains located on that agent's client machine throughout the query. The query is executed by providing the *query method* to the client machine, processing the method with local information, and then returning a response or taking an appropriate action (e.g. displaying product X for price Y). A central application of client-side data mining is to *privacy-protected customization*, where a vendor wishes to customize its products and services for a particular agent that is registered with SDI, and take advantage of personal information relating to that agent that is not generally available. As another application, we describe a client-side advertising auction, where advertisers compete for the right to display a banner advert to a user, based on local information about the user's preferences.

#### 4. Applications: Very Brief Overview

In the next section of the SDI description we describe some specific variations and systems that can be implemented within the general architecture. In overview, we describe the following key applications:

- **Safe user profiling and personalization.** This allows on-line users to receive personalized information and services without providing personal information to vendors, so that users retain control over their personal information. Users can interact with vendors under different pseudonyms, and provide information to the central SDI data warehouse to allow data mining. Finally, users can allow specific vendors to execute queries, where the result of the query is information that allows that vendor to customize its service.
- **Client-side user profiling.** The client-side SDI proxy can monitor the browsing behavior of a user, and submit data periodically to the central SDI data warehouse with appropriate meta information to provide the user with guarantees about the type of information that can be released to vendors during query execution.



- 1       • Static Data mining applications. Agents can submit queries to the data warehouse  
2       query-execution module and perform data mining and collaborative filtering on  
3       aggregated and anonymous information provided by data submitted to the  
4       warehouse by agents.
- 5       • Interactive data mining applications. Agents can request actions from the SDI  
6       system if certain conditions are found to exist in information, for example SDI  
7       can send information to other agents about services or products, if good matches  
8       are found. Another example is a request that a certain number of agents with  
9       particular properties be contacted and asked to take a particular action, we give an  
10      example within a transportation domain.
- 11      • Pro-active data mining. The system of SDI might itself pro-actively execute data  
12      mining queries, and index and classify certain types of data to allow more  
13      efficient future query execution, and to also suggest useful information to its  
14      client agents. The answers to popular queries can be priced and cached, so that  
15      providing agents receive value whenever an answer is sold to another agent. The  
16      system of SDI might advertise a set of queries to allow vendors to select pre-  
17      computed results.
- 18      • Matchmaking applications. We described in the top-level description of SDI a  
19      technique to implement “persistent queries”, which reside on the central data  
20      warehouse’s query execution module, and are triggered whenever the correct  
21      conditions exist in the data base. One typical use of such a query is to say  
22      “introduce me to other users with property P”, such that the system introduces  
23      user A\_1 with user A\_2 if A\_1 has the property required by A\_2 and also A\_2 has  
24      the property required by A\_1. This is equivalent to “introduction by mutual  
25      consent”, and is possible within SDI without *information leakage*, because the  
26      only agents that are informed of a match are the agents with the correct properties.  
27      Applications exist to finding a business partner, funding a new start-up (incubator  
28      Co.), forming an interest group, n-way negotiation, introducer system (by mutual  
29      consent).

## 1. Static Data Mining Applications

2

3 This section describes specific applications of SDI-based static data mining. A central  
4 example is collaborative filtering and personalization applications in electronic  
5 commerce, where consumers and vendors provide information to SDI, and the  
6 information can be queried within the price and data-access rules placed with the data by  
7 owners of the information. We refer to the variation of SDI with consumer profile  
8 information as the “iamworthit” system, because users can place profile information in a  
9 shared database and receive payments in return for queries performed by vendors.  
10 Vendors can use the profile information to build better customization models, and  
11 provide customized products to customers based on their profiles and what has worked  
12 with other customers with a similar profile.

13

We limit our attention in this section to “static” datamining, which as defined in the mid-level SDI description allows agents to query the data, but without contacting the agents to request more information. All queries are performed on the basis of the information already submitted to the database, and the rules associated with that information. Within the set of *static* queries, we do allow a vendor to identify a specific agent within a query command, so that if a user provides an identifier to a vendor then the vendor can query the database with its collaborative filtering model and determine an appropriate action to take.

22

It is possible to request an *action* as the result of a query, for example “send message X to all users with profile information Y”, so long as this message is not contingent on as yet unknown information about the user. We describe specific examples, for example to a smartbrowsing system, an education portal, and an advertising network.

27

One of the core purposes of SDI is to provide a common location and format for information that has been gathered from a wide variety of sources and that might require different sorts of analysis. Since its framework is designed to handle different types of data and algorithms, SDI can be used as a platform to explore and exploit the rich







many standard methods that can be used to achieve this, such as Principal Components Analysis. Another approach is to adjust the granularity of the data, if at all possible. In a music store analysis, for example, there might be many more album titles that artists (since each artist can produce multiple albums). In such a case, purchases could be recorded by artist rather than by album, greatly reducing the dimension of the customer vectors' purchase space.

### 1.1.1 Data Structure

In this application there are many types of information which can characterize both users and items. SDI is intended to function as the intermediary between a vast web of vendors, on the one hand, and individual consumers, on the other hand. Major sources of data include:

- 1) Demographic. Such data will most likely be elicited by SDI from vendors and consumers when they initially register for the service, and details very general characteristics about them. It will consist of numbers and categorical values (age, zip code, sex, level of education, etc.).
- 2) Commercial. This is the kind of data any that vendor collects in the course of doing business (especially e-commerce); generally, it links customer codes to purchase items, dates, quantities, and prices. Depending on the nature of the business, this data could be fairly complex, and might well include text. For example, one could imagine that a bookstore, in addition to keeping track of its sales history, collects book reviews, author profiles, and plot summaries.
- 3) Behavioral (vis-à-vis the Internet). A user's client-side SDI proxy can monitor his/her browsing behavior on the World Wide Web, monitoring the pages a user hits, the click stream and content requested, etc. Click stream information can be useful, for example, because it can indicate a user's interest in the information that it is presented with.

We assume that vendors and/or a third-party annotate web pages with *tags*, that provide a commentary of a web page and allow meaning to derived from a user's browsing behavior.



1 category  $i$  is represented as an  $n_i$  dimensional vector,  $y_i$ . Hence, the total number of  
2 dimensions used to describe the full set of  $n$  categories ( $y_1, \dots, y_n$ ) is

$$\sum_{i=1}^n n_i$$

3

4 Note that sparse methods are especially useful here, since a categorical vector  $y_i$  will  
5 typically consist of mostly zeroes, with a single non-zero coordinate representing the  
6 categories' value (i.e., we encode the color red, using the previous example, as  $(1,0,0)$  ).

7 Note also that category vectors with different values are treated as orthogonal by the system.

8 A final issue is the representation of text. As described in previous related patents, all  
9 relevant blocks of text in the database are converted into a dictionary that maps unique  
10 strings to the number of times they appear in the database. An appropriate TF/IDF weighting  
11 function is chosen and calculated for each of the  $p$  words that appear in the dictionary. The  
12 full set of text connected to a single customer can thus be represented as the vector  $(z_1, \dots,$   
13  $z_p)$ , where each  $z_i$  equals the number of times the word  $i$  appears in text related to the  
14 particular customer multiplied by the TF/IDF score assigned to word  $i$ .

15 In summary, when a database describes its customers using a combination of numerical  
16 values, categories, and text, customer  $i$  can be represented by the vector  $c_i=(x_1, \dots, x_n, y_1, \dots,$   
17  $y_m, z_1, \dots, z_p)$ .

### 18 **1.1.5 An Example Profile Vector**

19

20 Suppose we have a database containing information on customers' ages, their musical  
21 preferences (i.e. an answer to a survey asking: "Which do you prefer, Mozart or the  
22 Beatles?"), and the contents of the emails they've written. Furthermore, suppose the only  
23 salient variables in all the emails written consist of the words "Beatles", "Mozart", and  
24 "practice", and that we are using the function

$$TF / IDF(x) = \frac{1}{\sqrt{n_x}}$$

25

26 Where  $n_x$  represents the number of times word  $x$  appears in the dictionary. We now want to  
27 represent one of the customers in the database; he's a 10-year-old boy who prefers Mozart to  
28 the Beatles, and who wrote an email to his friend that mostly describes his attempts at



1 practicing Mozart, but in passing mentions his sister's new Beatles CD. Suppose he uses the  
2 word Mozart 2 times (although it appears 456 times in the full database of all customers'  
3 emails), the word Beatles 1 time (appears 217 times in database), and the word practice 3  
4 times (appears 77 times in database).

5 We define the following coordinates:

6  $x_1 = \text{age} = 10$

7  $y_1 = \{\text{Mozart, Beatles}\} = (1,0)$

8  $z_1 = \# \text{ of times customer uses word "Beatles"} \times \text{TF/IDF}(\text{"Beatles"}) = 1 * 0.067 = 0.067$

9  $z_2 = \# \text{ of times customer uses word "Mozart"} \times \text{TF/IDF}(\text{"Mozart"}) = 2 * 0.047 = 0.094$

10  $z_3 = \# \text{ of times customer uses word "practice"} \times \text{TF/IDF}(\text{"practice"}) = 3 * 0.114 = 0.342$

11

12 *In our example, then, we might encode this boy as customer 1:*

13  $c_1 = (x_1, y_1, z_1, z_2, z_3) = (10, 1, 0, 0.067, 0.094, 0.342)$

#### 14 **1.1.6 Choosing an Appropriate Level of Data Granularity**

15

16 We define the term granularity to denote the level of detail available within a given set of  
17 data, which is often structured hierarchically. Suppose a grocery store database contains  
18 records for a box of flavored gelatin powder. This could be categorized in a variety of ways;  
19 moving from the most specific to the most general, we might treat this data point as "12.5  
20 ounce, strawberry flavor, Jello-brand gelatin dessert" (which would be entirely different  
21 from "12.5 ounce, banana flavor, Jello-brand gelatin dessert"), or as "12.5 ounce Jello  
22 gelatin" (a categorization which would treat as identical the strawberry and banana Jellos),  
23 or as "flavored gelatin", or as "dessert", or as "food", or as "grocery".

24 When analysis is performed on such data, the level of granularity chosen will have a strong  
25 effect on the outcome of the analysis. If the level of granularity is too fine-grained, the data  
26 will be too sparse, although it could be potentially aggregated to the next highest level of  
27 granularity. If the granularity is too coarse, the results of the analysis might be overly  
28 general (e.g., a customer would find a collaborative filter useless if the only  
29 recommendation it makes for a dessert choice is "go to the grocery section of the store").

1 Since the level of granularity will have a salient effect on the outcome of an analysis, it  
2 should be chosen very carefully, and might well play a factor in pricing when a vendor  
3 chooses to sell its data.

#### 4 **1.1.7 Statistical Methods for Data Analysis**

5  
6 In order to perform a wide range of analytical tasks, SDI needs to make use of a variety of  
7 computational approaches. These are described below, starting with the simplest methods  
8 first.

##### 9 • (1). Standard Database Searches

10  
11 Since most of the data will be stored in centralized databases, simple searches,  
12 queries, and data filters can be implemented by means of standard SQL commands.  
13 Typically, data will be collected or sorted using efficient database calls before being  
14 fed through analysis routines; once complete, the results can be fed back out to the  
15 database environment for further efficient manipulation.

##### 16 • (2) Metrics – Measuring the Similarity Between Profile Vectors

17  
18 Given two customer (or vendor) profiles,  $c_i$  and  $c_j$ , it is frequently desirable to know  
19 how similar they are. For this purpose, we define the similarity metric  $M(c_i, c_j)$  to be a  
20 function that takes as input two customer vectors and returns as output a numerical  
21 value in the range  $[0,1]$ . When two customers  $c_i$  and  $c_j$  are identical,  $M(c_i, c_j)=1$ ; when  
22 they're completely different,  $M(c_i, c_j)=0$ .

23 The problem is somewhat simplified by the fact that we treat all customers as vectors.

$$M(A, B) = \cos \theta = \frac{A \cdot B}{\|A\| \cdot \|B\|}$$

24 Given two customer vectors, we can use the correlation between them to serve as our  
25 metric:

26 Note that  $\theta$  here represents the angle between the vectors  $A$  and  $B$ , and that we expect  
27 all coordinates of the vectors to be positive (in order for  $M(A,B)$  to keep its output in  
28 the range  $[0,1]$ ).

In more complicated cases, however, a customer vector might contain multiple fields with varying ranges of values. For example, we might have customer vectors of the form  $c_i = (\text{age}_i, \text{income}_i)$ , in which the maximum age is 80, but the maximum income is 300,000. In such cases, the coordinates with larger values will dominate the similarity metric, overwhelming any influence that smaller fields might have.

This requires a normalization of the customer vectors, which can be done in several different ways. One approach would be to scale every coordinate by the maximum observed value, forcing all coordinates to lie between 0 and 1 (again, enforcing the rule that all coordinates must be positive).

$$c_i = \left( \frac{\text{age}_i}{\max(\text{age})}, \frac{\text{income}_i}{\max(\text{income})} \right)$$

The only problem with this is that if a coordinate's maximum value is an outlier (being vastly bigger than the typical value), most of the coordinates' values will seem unusually small once they are scaled by the maximum. In such cases, it might be better to scale the values with a "squashing" function such as the sigmoid, which deadens the impact of extreme values; one such configuration would be the following:

$$\begin{aligned} \overline{\text{age}_i} &= \frac{\text{age}_i - \text{mean}(\text{age})}{\sigma_{\text{age}}} \\ \overline{\text{income}_i} &= \frac{\text{income}_i - \text{mean}(\text{income})}{\sigma_{\text{income}}} \\ c_i &= \left( \frac{e^{\overline{\text{age}_i}}}{1 + e^{\overline{\text{age}_i}}}, \frac{e^{\overline{\text{income}_i}}}{1 + e^{\overline{\text{income}_i}}} \right) \end{aligned}$$

Note that the mean and variance of the data points are used to fully normalize them, such that the sigmoid function will spread the values somewhat more evenly between zero and one.

The previous approaches are especially useful for single numerical fields, which might well overwhelm each other if some sort of normalization isn't performed.

1 A different problem arises for text or large categorical fields, since they can  
2 potentially consist of hundreds of coordinates capable of overwhelming the influence  
3 of single numerical fields. Suppose we believe the age of a customer is as important  
4 as the text of articles read. In such a situation, the thousands of coordinates devoted to  
5 the text field would dominate the metric's behavior, negating any influence that age  
6 would have on our measure of similarity – clearly not a good situation.

7 A solution to this would be to find the correlations among the fields taken separately,  
8 then average the result. That is, if each customer  $c_i = (age_i, text_i)$ , where  $text_i$  is a  
9 vector with a very high number of dimensions, we could define the metric:

$$M(c_i, c_j) = \left( \frac{corr(age_i, age_j) + corr(text_i, text_j)}{2} \right)$$

10 Where

$$corr(c_i, c_j) = \frac{c_i \cdot c_j}{\|c_i\| \cdot \|c_j\|}$$

11  
12  
13 The result is a metric that gives equal influence to each field.

14 • (3) Forming Vectors Into Groups

15  
16 The process of classification is essential to collaborative filtering, as it allows  
17 different vectors to be formed into groups based on some measure of similarity. If we  
18 are able to create groups of customer vectors, for example, we can then give  
19 individual customers recommendations based on the patterns of their group-mates,  
20 who presumably have similar tastes.

21 K-means Clustering and Nearest Neighbor algorithms are extremely useful for  
22 grouping purposes: previous iReactor patents give a full and detailed description of  
23 our customized versions. This section gives a brief overview of these methods.

24 (3.1) Clustering

25 K-means Clustering is an algorithm used to partition a coordinate space such that all  
26 vectors in a given partition are more similar to that partition's vector average (the

centroid), than to the centroids of any other partition. It is a process that iterates over the following steps:

0. "Seed" the coordinate space with the initial centroids, which are vectors used to describe the centers of the clusters, in the sense that they are the average of all the vectors currently assigned to the partition. This can be done randomly (assigning centroids random coordinates) if no other information is available, or it can be guided by pre-existing information. For example, if we wish to cluster vectors of music customers, we can use information about musical genres to create initial partitions that correspond to pop, gospel, classical, etc. This will locate the centroids in well-spaced intervals across the coordinate space.

1. Assign vectors to the most similar centroids. This is done for each vector by scanning across all centroids and calculating similarity  $M(\text{vector}, \text{centroid}_i)$ ; once finished, the vector is assigned to the cluster whose centroid has the greatest similarity. In this stage, vectors may switch their allegiance from one centroid to another, if the relative distances to the vector have changed sufficiently since the previous iteration. If no vectors change their allegiance, the iteration process is complete, and the algorithm stops.

2. If the iteration is not complete, recalculate the centroids by setting them equal to the average of those vectors that have been assigned to them. Go back to step 1.

Once the algorithm converges, the vectors are grouped into clusters. The centroids' coordinates as well as the identity of cluster members is useful information that can be passed on to subsequent stages of analysis.

### (3.2) Nearest Neighbor

The nearest neighbor algorithm, simply stated, creates a list of those vectors in a database that most resemble a particular target vector. This is accomplished by comparing the target vector, in turn, to every other vector in the database; the similarity between them is recorded, and once the comparison loop is complete the list of similarities is sorted. The top k members of this list are returned as representing those k vectors which most resemble the target.

- (4) Generalizing Across Databases



1 Suppose we have a set of databases, A, B, ..., Z. Taking each database in turn, we  
2 cluster it using all available data. Thus, using every record in database A, we group  
3 A's customers into clusters  
4  $A_1, A_2, \dots, A_n$ . Taking database B, we create clusters using all of B's information,  
5 creating customer clusters  $B_1, B_2, \dots, B_m$ , and so forth.  
6 Now, scan both databases for common pseudonyms (representing those customers  
7 who have interacted with both vendors under the same pseudonym) and create count  
8 variables  $w_{ij}$  to represent the number of pseudonyms that appear jointly in  $A_i$  and  $B_j$ .  
9 We can now produce the probability that a pseudonym appearing in  $A_i$  will appear in  
10  $B_j$ :

$$P(B_j | A_i) = \frac{P(B_j \wedge A_i)}{P(A_i)} = \frac{w_{ij} / total}{\sum_{j=1}^m w_{ij} / total}$$

$$total = \sum_{i=1}^n \sum_{j=1}^m w_{ij}$$

11 For example, if we have a database of airline ticket purchases and a database of  
12 restaurant visits, we can create clusters, in the first case, of customers who travel to  
13 similar destinations, and in the second case, of customers who eat at similar  
14 restaurants. Given that a particular customer belongs to a cluster of people who  
15 frequent Caribbean restaurants, we can infer which travel packages would most  
16 appeal to him based on the linking probabilities, as defined above.

17 • Multivariate Extensions:

18  
19 If we have a third database C, and there are a large number of pseudonyms common  
20 to A, B, C, the above probabilities can easily be extended. For example, knowing that  
21 a customer appears in  $A_i$  and  $B_j$ , we can calculate the linking probabilities to any  $C_k$ :

22





counted, for TF/IDF purposes, by accompanying dictionaries), the union of the words will define the text coordinates of the new common-information profile. When word counts are being mapped from their original databases to the new vector, the original TF/IDF weightings may be used, or new TF/IDF weightings may be created (using a dictionary constructed from all the databases' text taken together).

Once analysis has been performed, certain common-information profiles will be grouped together by their shared similarities, although the pseudonyms they represent may have been originally drawn from different databases. Such groups will represent links between different databases, and may be used for predictive purposes (see end of example).

### **1.1.8 Example of Cross-database Analysis**

In this example, suppose that the central SDI data warehouse contains data submitted by the following vendors:

- A. A travel agency keeps track of tickets sold, and vacation web pages browsed.
- B. A bookstore keeps track of books sold, and stores an electronic version of the New York Times Review of Books.
- C. A sporting-goods and clothing shop, keeps track of purchase items sold (which includes magazines, for which electronic text exists).

A certain airline wants to promote various vacation packages it has available, which include both European and Caribbean vacations, as well as singles and family packages. Although it has leased rights to databases A,B, and C, it turns out that no customer pseudonyms appear in more than one database at a time – in other words, there are no shared records.

A vacation expert is hired to create a common-information profile. He creates the following information vector: (list of tropical countries, list of European countries, family score, list of sports, text)

Note that the family score is a numerical value ranging from 1 (young singles) to 10 (many small children), and indicates what kind of person the customer is (a party-oriented student vs. a sedate father of three).

The expert the creates the following mappings:

- 1       A. Travel Agency. Link destinations of tickets sold to country fields (i.e., the number
- 2       of trips to Germany by a customer would be placed in the Germany field of the
- 3       common-information profile). Link sales of children's tickets, or requests for
- 4       children's meals, to family score. Put web-page data into text field.
- 5       B. Bookstore. Link travel books' text to country lists. For all books purchased by a
- 6       customer, map text from book reviews into text field.
- 7       C. Sporting-Goods store. Map warm-weather clothing (and swim gear) to tropical
- 8       countries, ski gear to countries with skiing areas. Map sales of toys or children's
- 9       clothing to high-value family scores, map revealing-bikini and student-discount
- 10      sales to low-value family scores. Map text from magazines purchased by a
- 11      customer to text field.

12  
13   These mappings are then applied to each database, generating a full set of common  
14   information profiles. These are then clustered, forming groups that share commonalities.

15   The expert can now do several things with the results. First of all, he identifies the general  
16   "flavor" of each cluster (e.g., families with small children that enjoy winter, Europe, and  
17   skiing); the pseudonyms contained within each cluster can then be targeted for vacation  
18   packages suitable to their tastes. Secondly, the fact that pseudonyms from different  
19   databases have been clustered together allows the expert to plan cross-category marketing. If  
20   certain travel-book-buying parents have been grouped together with parents who bought  
21   their children swimsuits and scuba toys, it may be that they share a preference for family  
22   activities that take place in warm places or by the seashore. Hence, the book-users might be  
23   advertised various ocean-related sports goods appropriate for young families, and likewise  
24   the swimsuit-users might enjoy getting recommendations for travel books that describe  
25   tropical destinations that are especially fun for children. That is, if the goal is to cross-  
26   market items from A to customers in C, the most logical source of recommendations would  
27   be the people in A who have been grouped with the people in C.

### 28   **1.1.9 Methods for Validation**

29  
30   To a large degree, the overall success of an SDI analysis is the relevance of the connections  
31   that are inferred from the data. It is often the case that a certain amount of validation is



adding it to a running total. The system with the highest total can thus be judged the most effective, since it most strongly recommended items that the customers did, in fact, end up purchasing.

Because the result of this type of validation is a quantitative score, it is possible to automate the model selection process. Given a set of analytical approaches (each with its own array of parameter settings), it is possible to loop through the full parameter space (using a grid of evenly spaced numerical values, if needed, to reduce dimensionality), computing a validation score at each iteration. Those combinations of algorithms and parameter settings that demonstrate the best performance could be chosen as the top candidates for the final system configuration, since they do the best job at predicting customer behaviors.

#### (1b) Quantitative Approaches – Dynamic Method

The problem with the hold-out approach to validation is that it isn't dynamic, since it doesn't reflect the impact that the recommendation system has on the customers once it is implemented, and may be based on data that doesn't contain current trends. After all, it is better to predict what the customer will buy rather than what the customer has bought in the past.

A better approach is to run a controlled experiment against the actual customer base. First, the pool of customers is split at random into different segments. Next, each approach under consideration is used exclusively to make predictions for a given segment. Once the trial period is over, each system is given a score based on how valuable its recommendations turned out to be (this could be measured by total sales generated, for example, or by the number of times a customer made use of a recommendation).

#### (2) Human Expert in the Loop

Although quantitative methods can automate the validation process to some degree, at the beginning of many projects there is so much raw input data available and so many decisions that have to be made about the analytical approach that an automated process would have to test a prohibitive number of combinations of data, algorithms, and parameter settings to get optimal results. In such cases, it is useful to employ a



1 Iamworthit allows individuals to receive payments for the information that is collected by an  
2 SDI client proxy and stored in the central SDI data warehouse, for data mining purposes.  
3 Iamworthit also allows individuals to receive payments in one-to-one interactions with  
4 vendors in return for providing vendors with information that allows them to make an  
5 appropriate offer. Secure data interchange supports the useful exchange of information  
6 between agents without allowing vendors to collect and distribute information about users  
7 without the knowledge or permission of users, as is possible now via the system of cookies  
8 and affiliate networks, such as that operated by [www.doubleclick.com](http://www.doubleclick.com). In the system of  
9 Iamworthit vendors benefit through well-targeted advertising (both push and pull), the  
10 ability to customize information and services (even to first-time customers), and access to a  
11 large database of information about buyer purchasing habits. Individuals still release profile  
12 information to the central SDI database for the purposes of controlled access by vendors,  
13 and also provide vendors with profile information during interactions and allow vendors to  
14 execute queries on information and receive the benefits of personalization without directly  
15 accessing the information. Users can maintain multiple pseudonyms and profiles, but within  
16 SDI vendors can still access information across pseudonyms and use cross-web (broad) and  
17 single-vendor (deep) information to build robust models of buyer behavior. Buyers benefit  
18 through personalization with privacy, and financial rewards in return for releasing profile  
19 information.

20 Within B2C e-commerce, secure data interchange can also support a system of *time-of-*  
21 *purchase competition*, which allows an individual to use the profile management capability  
22 of SDI to provide vendors with information about a user at the time of purchase, and allow  
23 competing vendors to offer the same product at a better price, or a better product (for the  
24 user) at a good price. With time of purchase small entrants to the marketplace can make  
25 counteroffers to users that are about to purchase a product or service from another vendor,  
26 and can compete in small parts of the marketplace without investing heavily in advertising  
27 and brand awareness. Time-of-purchase requests can be made by user clients to iamworthit,  
28 to request that iamworthit cascades purchase requests onto other vendors, collects responses,  
29 and then return them to user clients. This extension of SDI is discussed in Section 2.1.

30 As a commercial strategy, one might make it a necessary condition of belonging to  
31 Iamworthit that client machines submit profile information to the SDI central data

warehouse, so that vendors can perform useful data mining and then provide customized products and information to users, for example based on models of collaborative filtering.

### 1.3.1 System Overview

Client machines in Iamworthit implement an SDI proxy on-top of a user's regular web browser. The role of the proxy is:

- a) Profile and identity management as the individual interacts with server computers of on-line vendors.
- b) Collect and manage profile information, with information periodically submitted to the central SDI data warehouse.
- c) Release profile information to vendors according to profile management policies.

In Iamworthit the client-side SDI proxy is configured by users to periodically push information collected about the user to the central SDI data server, i.e. click stream data, profile information, purchases made, information requested, etc. Profile information is associated with price-rules, as in the top-level description of SDI, and can also be adjusted with random perturbations to protect the identity of a user under multiple pseudonyms.

The client periodically sends update messages to the database, for example with information about new activity (e.g. web browsing, purchases, adverts not selected, etc.) SDI allows the information to be used by vendors to build models that enable good personalization of products, services and adverts.

Users might choose to randomize profile information to protect sensitive information and to prevent vendors linking users across different pseudonyms. Clients submit perturbed data points and the distribution that was used to generate the noise (see the section on random noise perturbation).

### 1.3.2 Privacy-protected data mining

Vendors can gain access to profile information in the central SDI data warehouse, as long as they meet the criteria of the agents that submit information on behalf of users, and pay the price of data access. The profile information can allow vendors to build complex models for personalization, negotiation, and advertising, based on information about purchases made by different types of users.

1 Users submit profile information to a central database that can be used for profiling, without  
2 revealing their identities. Client-side SDI proxies track user activity on-line, across multiple  
3 sites, and submits data, randomized if necessary, and only according to a user's preferences  
4 to the central SDI data server. The data is useful for building models of buyers, for example  
5 purchasing patterns, for the purposes of personalization of information and adverts. Agents  
6 that submit information retain ownership of the data.

7 Simple data mining queries include:

8 (a) Compute the average income level of people purchasing camcorders. Suppose  
9 that a vendor has pseudonymous identities of its recent customers, and wants to  
10 compute their average salary. It does not know the salary of any of the customers,  
11 and individual customers will not release their salary to the vendor. However, if  
12 some of those customers have provided information about their salary to the  
13 central SDI data warehouse then the vendor can compute the result to its query.  
14 Agents might associate price rules with information about their salary that allow a  
15 vendor to query that information so long as the identity of the agent is not  
16 revealed, i.e. so long as the salary is provided anonymously. We described earlier  
17 in the patent description how price rules can define different prices for different  
18 types of information access.

19  
20 (b) Compute the total donation to a fund for bone cancer research. Individuals do not  
21 wish to release information about their specific charitable donations, but might be  
22 happy to release that information anonymously. Again, a query to compute an  
23 average donation can be formulated and executed in the central SDI data  
24 warehouse.

25  
26  
27 (c) Perform collaborative filtering across multiple fields. This is explained in more  
28 detail below, and is possible within SDI because a general query can be  
29 performed so long as the querying agent makes payments for data access in  
30 accordance with the price rules of agents which submit information to the shared  
31 database. Note also that randomly perturbed data does not prevent the



computation of correlations between fields, so long as the random perturbation is “small” with respect to the value of the data record.

### 1.3.3 Importing Off-line Data into the Data warehouse

The central SDI server can associate off-line information about a user with a user’s on line pseudonymous profile, even though the central server does not know the user’s pseudonym IDs. This can only be done with the user’s consent, and may also involve appropriate compensation. Within the system of iamworthit we can credit users for both off-line and on-line information.

Merging a marketing database with SDI user profiles can be useful both to initialize the database, for example when asking a user questions to generate an accurate user profile rapidly and efficiently. Off-line data can also add useful richness to on-line profiling information, which may be largely contextual and low on details/factual information. For example, off-line data can include information such as whether a user owns a car, rents an apartment, has house insurance, life insurance etc. SDI can also extrapolate correlations to other user profiles, on the basis of common SDI-profiles, for example using statistical techniques.

It is often the case that individual customers appear in some databases, but not in others. Under normal circumstances, an analyst working across different databases would be faced with a large number of incomplete customer records, each with gaps corresponding to the fields of the databases to which they don’t belong. A solution to this problem is offered by SDI, which is capable of drawing correlations between different databases - this information can be used to generate predictions to fill in the gaps of incomplete customer records. The result is a full set of customer records that can be meaningfully sorted or filtered by any of the combined fields, and which can now be handled as a unified set of data, suitable for use by standard database analysis systems.

In a typical example, SDI might be used to combine a demographic database, such as the one offered by the Econometrics Corporation, with a commercial database, such as the one offered by Claritas. The Econometrics database consists of 180 million different customer records, but at a fairly coarse-grained level of detail, consisting of such information as age, gender, family status, location (at the state, city, or zip code level), and personal income. In

1 comparison, Claritas offers a smaller base of customers, but includes information of  
2 arguably higher quality, since it breaks customers down to the geocode (sub-neighborhood)  
3 level, and includes much more detailed information on personal spending habits across  
4 hundreds of different purchase categories. A logical reason to combine these databases  
5 would be to supplement information about customers in the vastly broader demographics  
6 dataset with particular predictions about their personal preferences and likely commercial  
7 spending habits. One could imagine using this augmented data set to support a web site that  
8 instantly customizes itself to new visitors' preferences. Since the number of records in the  
9 Econometrics database is equivalent to roughly 72% of the population of the United States,  
10 it is likely that most first-time visitors to the site will already have a "thumbnail sketch" in  
11 the system, and can thus be greeted with an page appropriately configured to their personal  
12 tastes.

13 The technical details of the combination process (which have been described elsewhere in  
14 the patent) to a large degree depend on the amount of overlap between the databases, that is,  
15 the number of customer records which are shared in common.

16 Suppose the demographic databases' fields are coded  $(x_1, \dots, x_n)$ , and the commercial  
17 databases' fields are coded  $(y_1, \dots, y_n)$ . Suppose further that customers in set A appear only in  
18 the demographic database, customers in set B appear only in the commercial database, and  
19 customers in set C appear in both.

20 The process of supplementing the fields of customers A depends completely on the  
21 derivation of the distribution  $f(y_1, \dots, y_n \mid x_1, \dots, x_n)$ , which describes the correlation of fields  
22 in the commercial database on fields in the demographic database. As previously discussed  
23 in the patent, different techniques may be used to create this distribution, depending on the  
24 size and variety of C.

25 As a concrete example, one could imagine that set C includes customers from rural areas.  
26 The demographic database would reveal that, although their incomes aren't huge relative to  
27 the national average, they tend to spend a lot of it (i.e. are active consumers), have large  
28 families, and purchase large vehicles. The commercial database might show that they  
29 enjoying hunting magazines and Ford trucks. If they live inland, they buy hunting  
30 equipment, if they live near the ocean, fishing equipment.

1 If these trends are dominant in set C, they will impact the distribution function. Thus, when  
2 a browser from a small town in Texas with a typical income pattern visits the automated  
3 website, he could be greeted with discounts on truck accessories and a small sidebar with  
4 news on the hunting season. On the other hand, a visitor from a small town in Maine might  
5 be given the same truck discounts, but would have news on the fishing season.

6 Although the demographic dataset is arguably the weaker of the two in terms of content, the  
7 fact that it contains even a small amount of information on most people in America makes it  
8 very valuable for handling first-time visitors, since most of them will appear in it. By using  
9 SDI to leverage the more detailed information in the commercial database, we are able to  
10 supplement the rough demographic data with predicted commercial preferences. This allows  
11 us to construct more detailed thumbnail sketches for each customer, allowing our reception  
12 of first-time visitors to be much more appropriate (since knowing personal hobbies or  
13 interests tells us much more about a person than general income level).

### 14 1.3 Static Query Execution: Central Applications

15  
16 Some key applications of SDI are:

- 17  
18 • Assessing the Value of Data. Plug together sets of data, and measure predictive  
19 accuracy.
- 20  
21 • Matching Data Across Vendors. Find patterns in common pseudonyms, denoting  
22 common areas of interest; use catalogues of order codes and item description to  
23 find similarities across data sets.
- 24  
25 • Targeted Recommendations; e.g. match customers to their nearest neighbors in a  
26 data set and generate recommendations for users, collaborative-filtering style  
27 application.
- 28  
29 • Leveraging Portal Data. Use data from portal to leverage data needs for ISP  
30

- 1. Analyzing Affinities. Suppose a vendor has a list of customers, and knows to some degree what web pages they visited after leaving vendor site. A large collection of customers taken from an ISP will contain their web-surfing behavior. Cluster web sites and cluster customers, finding cluster-to-cluster interactions. Use this information to classify vendor's customers; gives vendor an edge in knowing customers' tastes.

8 There are a number of search-based applications, where SDI searches for appropriate  
9 profiles and then requests that SDI makes contact with the users pseudonymously, i.e.  
10 without the vendor receiving any useful information about a user's identity. The contact,  
11 interaction, and business relationship with the vendor occurs under terms of complete buyer  
12 pseudonymity. In accordance with the parent patent application [INSERT US PATENT  
13 No.] the pseudonymous communication may be either email, real-time text  
14 communications, voice (such as the pseudonymous telephony or Internet telephony). In the  
15 case of pseudonymous telephony, instead of a one-time or persistent pseudonymous buyer  
16 address, pseudonymous buyer telephone numbers may be used for the third party to reach  
17 the buyer under his/her terms. Example applications include:

(i) Financial Advice and Financial Planning Services. Often buyers are quite sensitive about the confidentiality of the release of this type of information related to personal financial matters and particularly with certain matters (and perhaps in general) prefer that their financial advisors were unaware of their true identities. Similarly, investment advice or sales communications by stock brokers are another application where similar buyer information is typically disclosed.

(ii) Insurance Agents & Brokers. For many types of insurance, (e.g. health, life, casualty) personally sensitive information is disclosed by buyers to their agents and brokers. Initially, before insurance services are purchased, it is possible that useful detailed quotes and/or insurance advice could be provided to a buyer pseudonymously.

(iii) Legal advisors. There are a variety of legal disciplines in which the associated legal services delve into highly sensitive personal information (e. g., bankruptcy law, divorce law, criminal law, etc.) Many lawyers also offer to first-time prospective clients a free consult in which such a privacy-enhanced communications system could be initially beneficial to the parties.

(iv) Family Counseling and Psychological Counseling. The parent patent application also suggests these applications which often involve the exchange of highly confidential personal information.

(v) Medical Consultations Involving Drug Prescriptions. In this variation, although medical consultation may be conducted anonymously, in order for a physician to prescribe medication, the identity of the patient must be known (within the current regulatory legal requirements).

(vi) Advertising network (where the clustering or nearest neighbor algorithm interact with the ad server). Ads on that ad server's database, e.g., on the sites most frequently visited by the user are periodically uploaded to SDI such that the ad(s) of highest predicted interest to the user is presented upon the user visiting that site. Preferably the target object profiles of all the ads on the ad server database, as well as the profiling algorithm, which is used by the ad server is properly integrated with that of SDI so that it is possible to convert the usage statistics of the data model used by the ad network (which may include ad server data for users collected across the ad networks) into useful statistics for SDI.

(vii) Content sites, e.g., for personalizing news articles (which are again uploaded to the SDI server each news day if those articles would not otherwise be available on that web site. If they are, as with any web content format which is universally accessible, the content may be frequently profiled following retrieval, e.g., by a simple web crawler script.

(ix) **Rating Sites.** The ISP level proxy may also contain metadata relating to specific sites, products or informational content (as an alternative to the metadata residing on the HTML pages, which they describe through cooperation with the vendor). Such metadata includes annotations, average user ratings, according to a variety of attributes, as well as the ability to browse Web pages with associated annotations and/or user ratings provided from those users who are identified by either: a user-selected cluster, or the user's profile or high or low ratings according to site selected. Endorsements (and ratings by a variety of relevant criteria) by a variety of types of organizations may be available as well, and browsing and searching may be performed with these organizational endorsements/ratings as a criteria to bias or filter searches or filter the sites accessible via browsing interface (the same may be also performed for user based ratings as well).

- (x) **Cache Engines** – As an alternative to the profile generation and processing modules running on the network vendor (ISP) servers, it may operate instead,

1 or in combination as a distributed process upon the caching server in  
2 accordance with the method taught by "Broadcast System for Reduced  
3 Memory Devices and Asymmetric Networks". (check exact title of this  
4 patent) As taught in this patent (in accordance with the current applications)  
5 the pages predicted to be accessed by the user on the following day are  
6 precached in advance.

7  
8 Typically sites which tend to be visited most frequently are assessed for  
9 selections which are personally relevant but have never been previously  
10 accessed. These would be presented as recommendations to the user if/when  
11 he/she visits that site (which is also probabilistically determined). The fact  
12 that the recommendation (whether generated by the site itself or modified at  
13 the client level proxy), the user's behavior is positively being reinforced by  
14 the fact that the ranking of personal recommendations is prioritized in direct  
15 relation with the prioritization scheme for precaching, i.e. personalized  
16 recommendations and precached pages are substantially identical and as a  
17 result of precaching to the server, or even the client, can be accessed by the  
18 user with little or no latency (thus in the ideal embodiment it may even be  
19 advantageous to highlight the links which have been precached).

20 In a preferred commercial model, the techniques which are above described,  
21 that meta data containing profile information on the pages are encrypted such  
22 that a cache engine that is used to precache cannot be decrypted and read by  
23 certain caching engines which do not possess a desired business relationship  
24 with iamworthit. In another implementation, other competitive intermediary  
25 services to iamworthit may not be able to decrypt and read these web page  
26 profiles in order to present personalized information to the user (even if the  
27 competing protocol were pursued, the accuracy of the iamworthit version  
28 would be substantially more accurate due to its ability to leverage vendor  
29 centric SDI data from the host level proxy.

30  
31 (xi) Auto Insurance Application





(xiii) Medical information, such as medical conditions, medical history, active prescriptions, drug reactions, family history, possibly even genetic predispositions (from a genetic profile). Medical insurance information may also be potentially useful for a prospective qualified accessory to be able to readily access in case of an emergency.

(xiv) Physical location information—Users or advertisers could, for example: a) Query a pseudonymous user database to access profiles that are in close physical proximity and match certain criteria, e.g. live in a certain geographical region, had recently attended a meeting or event (or is planning to attend a particular event) had recently communicated with a friend or associate. In another variation, a user could for example, submit a query pertaining to every user in a particular physical space, e. g., a room, hotel or convention center, e. g., identify all users present here who attended Internet World, 1995.

### 1.3.1 Buyer Infomediary

The central SDI data warehouse can be used by vendors that provide services of *buyer infomediaries*, for example providing buyers in B2C e-commerce applications with historical information about previous purchases of users. At present companies such as [www.priceline.com](http://www.priceline.com) make profits because many individuals post take-it-or-leave-it buy offers for goods that are above the reservation price that vendors are prepared to accept for a service. An information infomediary, based around information submitted by users about previous successful or unsuccessful bids would provide for a more efficient marketplace.







1 experimented with such that optimal click-through or transactions occur particularly for key  
2 pages or purchasables. Adding to this statistical data SDI user profile information is able  
3 to reveal (via data mining) much more robust relationships between pages as these  
4 relationships change with user profile features. Accordingly appropriate rules may be  
5 provided for individuals based upon this general user profile information which is gleaned  
6 from interests and behavior before visiting the site, i.e., either from the user-centric or  
7 vendor-centric SDI (while the user visits the vendor's site or alternatively accesses an  
8 affiliate site).

9 In another variation of SDI we can highlight content on web pages with information that  
10 might be relevant to a user, even when the source of the web page is not personalized. This  
11 can be done via collaborative filtering techniques, which might bring in feedback and  
12 comments from other similar users within SDI that are stored in the central SDI database.

13 The iamworthit (user side) SDI database in cooperation with the vendor centric SDI  
14 service can sell to the vendor centric SDI service or other industry or market research  
15 organizations strategic information about the comprehensive behavior activities and user  
16 profiles of visitors and customers of these vendors (as is suggested earlier in the spec).  
17 Additionally, targeted survey questions may be presented on behalf of these entities in  
18 order to extract further information which may be correlated with certain features and  
19 attributes of these users. One such method, rapid profiling is detailed in the parent  
20 application [INSERT US PATENT NO].

21  
22 Vendors can utilize SDI's data mining interface to observe content and product  
23 consumption affinities based on user profiles describing Web wide behavior. This  
24 interface may also enable vendors to observe comparative click through and/or  
25 transaction rates of their competitors, how these criteria are affected by user profile  
26 attributes including geographic criteria (if relevant to that vendor) as well as how these  
27 criteria, (including even individual multivendor customers) are affected by various types  
28 of recommender and rule based engines and further the particular rules and weighted  
29 feature correlations used in generating these recommendations (again as described above,  
30 these rules and feature relationships may be derived automatically through a core sample  
31 of users, who are iamworthit subscribers). A vendor can use statistical data to tune



1 provider can query the data interchange for information about the user. The Internet content  
2 provider sells well-directed advertisements to vendors. Secondly, the data interchange could  
3 sell or rent data to an advertising agency directly, providing information in real-time to  
4 enable the advertising agency to provide more focus in its banner ads for its clients. "Per-  
5 transaction" pricing is a very powerful pricing model that is enabled with on-line banner  
6 ads. It is simple to monitor the number of click-through that are received at a particular  
7 banner, in response to an advertisement. In the off-line world pricing must be based on the  
8 number of impressions, or worst still, the number of mailings sent and it is more critical to  
9 understand the expected value of a campaign up front.

10 The proxy server could also act as an "ad network" itself, and sell focused advertisements  
11 for vendors, and purchase ad-space on the sites of content providers. The on-line domain  
12 provides this unique opportunity for quick experimentation with advertising strategies in  
13 order to get feedback on the likely utility of untested approaches. The system can use a  
14 hierarchical cluster tree to identify the most revealing items in a dynamically responsive  
15 fashion such that the profiles of all of the selections can be generated with the most minimal  
16 amount of interactions with the user (see "Rapid Profiling" section in issued patent entitled  
17 "System & Method for Customized Electronic Identification of Desirable Objects). Thus a  
18 more robust statistical model across multiple vendors is established as a result of the user's  
19 click through response of these intelligently selected virtual banners as well as other pages  
20 which are subsequently navigated through once the remote site is accessed via the banner.

21 In the preferred approach rapid profiling not only dynamically identifies and presents items  
22 which are most revealing of the other items in the collection, it also selects the users whose  
23 profiles suggest the greatest familiarity with these items (i.e., potentially correlated items).  
24 Furthermore, if the system's objective is to find new users or users who may be interested in  
25 the present vendor's other products, products for which little is known, then it will match  
26 users who are least familiar with exemplar items. The idea is to reveal the most significant  
27 data about the user profile with respect to the present collection of items of interest. Finally,  
28 rapid profiling can use direct explicit queries to determine interest on an item(s) or to collect  
29 demographic data on a user.

30 The target object profiles of advertisements on the ad server are matched against the user  
31 profile in order to automatically present the most relevant recommendation(s). Typically,





1 the preferred embodiment, any inserted ads delivered by iamworthit can be “turned off” by  
2 the buyer voluntarily unless the in kind value which the buyer receives in exchange for the  
3 advertising mandates a certain minimum ad impression delivery. Also in light of the above  
4 buyer concerns in the preferred embodiment, there are no additional impressions added to  
5 the buyer’s web browsing experience.

6 Ad blocking technology can be deployed to block existing banners and replace them with  
7 iamworthit ads. Ad blocking is commonly available, and the techniques used are well  
8 known in the art. The HTML source for the ad banners may be either called-up from the  
9 remote server (such as with an ad network) or alternatively from an ad server (typically  
10 purchased by the vendor) on the vendor’s own host server. In either case, the ad blocking  
11 software typically recognize the HTML source which originates from the ad server. A  
12 directory of the HTML source for the various ad servers is maintained with the ad blocking  
13 software (and presumably updated e.g. if new servers are added or change their IP  
14 addresses). The location/physical dimensions of the portion of the page occupied by the  
15 HTML source which the software removes (its “footprint”), may be readily replaced by  
16 another replacement banner (by iamworthit) from wherever the proxy server resides  
17 (typically on the client but potentially on the network). As a result of a potential difficulty  
18 in which the ad server deliberately changes its IP address to avoid recognition, which occurs  
19 constantly and dynamically, it may be possible to recognize portions of the page which  
20 contain image “features” of ad banners, which are inserted from another HTML source  
21 where that HTML source is linked to known vendor sites (which are also linked similarly  
22 from other site’s ad servers and/or are known via their HTML sources to have previously  
23 utilized ad server technology and a “new” HTML source appears instead, etc.

24 This model may be extended to other media domains e.g. replacing digital TV commercials  
25 instead with targeted ads delivered and precached via iamworthit for insertion at appropriate  
26 times i.e. during commercial breaks during standard video programming. Another variation  
27 could be tailored to pre-loading iamworthit advertising to automobiles for insertion in place  
28 of traditional radio commercials. In accordance with the parent patent application, (and as  
29 suggested above), email may be a useful targeted ad delivery medium as well. Per the  
30 patent case, the system is able to classify email (according to its source and content using  
31 implicit or explicit actions of the buyer). Based upon certain desirable confidence threshold

settings, the system may automatically delete “span” and replace it with targeted messages provided by iamworthit which the buyer is paid to receive pseudonymous physical mail is yet another potential source of targeted mail for which the buyer may be paid in conjunction with the user’s voluntary removal from direct mail marketing services which iamworthit could provide in conjunction. iamworthit’s buyer compensated targeted advertising model could be extended to a pseudonymous telemarketing service (as suggested above) using one-time call session pseudonymous telephone numbers for telemarketers to reach desired buyers. Finally, human sales persons could be the “ad delivery medium” for iamworthit.

### 1.3.6 Example: Dynamic personalization of links within a web site

Wide user profile data can also enable the dynamic personalization of links within a web site, to create a virtual shop floor to match the predicted preferences of each user. For a new user, that has never before visited the site, it is very advantageous for the site to already know about the preferences of that user in order to personalize the goods and services that it offers. The information provided at the secure data interchange, and gathered from the transactions of a user with another vendor, is vital for this type of personalization to first-time users. For return customers, an Internet site can also leverage its information that it has collected from previous interactions with the user, information that is collected locally to the site.

### 1.3.7 Personalized Search Engines

The present state of the art for search systems involves the use of an extremely static interface, which is not personalized for a user. In the previous patented disclosure [INSERT U.S. PATENT NO.] we identify features of a user and anticipate areas of content likely to be of interest. A personalized portal interface can then allow a user to view categories that are presorted in terms of expected value to a user, *based around more than the search term just entered.*

For example, if AltaVista knows that I am a research student, then the search engine can bias results in the direction of information that is likely to be appropriate to my interests, and away from other items. The personalization and filtering can be done based on results that I have accessed in the past, and also on the basis of information that is stored about a



1 following description), user profile which is “most like” the profile of the cluster to which it  
2 belongs, perhaps a median metric.

3 The Platform for Privacy Preferences (P3P) [CR 98; RC 99] provides for the ability to  
4 utilize XML meta-tags to annotate Web pages, and within a system such as SDI we can  
5 allow comments from previous visitors to a page to add annotations. Users can receive  
6 annotations from the SDI data warehouse that are associated with annotations provided by  
7 other users with similar profiles. One of the divisional applications of the parent case  
8 “System for Customized Electronic Identification of Desirable Objects” [Herz 98b] relating  
9 to the automatic creation of virtual communities suggests that users may be automatically  
10 assigned to particular communities (e.g. chat groups, forums, etc.) for this purpose.

11 We might allow users to rate the annotations in the pages, allow SDI to learn how useful  
12 annotations provided by certain users are to new users, and classify users as “experts” that  
13 receive priority in the position of their current and future comments. Future comments from  
14 users with a poor rating history for a particular content cluster may be deleted. A persistent  
15 interface feature on the tool bar or side bar can provide for annotations to also be accessed  
16 by users selecting certain profile features of users as they browse from page to page; for  
17 example a user could identify the comments of a news article about abortion by users who  
18 are self identified as advocates of the Women’s Rights Movement, ultra conservative senior  
19 citizens, teen women or those with a strong interest in alternative medicine or the Catholic  
20 Church.

21 The parent case [Herz 98b] further suggests that users may actively provide ratings in a  
22 completely privacy protected manner according to various criteria of pages they browse.  
23 Profile-based clustering of these ratings and annotations can allow a user to submit as a  
24 query a user profile, to receive a page that rates or annotates a page, and a listing of highly  
25 relevant pages to a particular type of user. This could also be used in a “reverse engineering”  
26 sense, for example requesting the exemplar user profiles for users that visited and gave high  
27 ratings to a particular site, and then using those profiles to find similar sites of interest, or  
28 finding the user features that are in the exemplar profile to understand something about the  
29 content of a web page. The browser interface may automatically display the related links  
30 that are determined dynamically to be most relevant to a user, again statistically estimated

1 via the data in SDI of users' browsing habits, and/or via active page recommendations or as  
2 book marks by those users as being of particular relevance or similarity to the present page.  
3 Browsing methods may include search-based browsing, and also browsing via a hierarchical  
4 navigation menu system, with users classified according to their behavior patterns and/or  
5 ratings which have been actively submitted. The parent case [Herz 98b] also suggests the  
6 use of hierarchical clustering for products, to help in a smart shopping system. Similar  
7 products can be presented together where the feature criteria for creation of the hierarchical  
8 cluster tree could be price or other criteria. In addition to allowing users to view item  
9 selections according to desired selection criteria, a comparison shopping function can also  
10 allow a user to view the attributes of buyers that tend to buy certain items to help to add  
11 confidence to a user that he/she is getting the right product.

12 The above description also describes the use of a hierarchical menu through which groups of  
13 users may be identified by their profile features (wherein a profile feature could even be a  
14 rating criteria itself of for example, an opinion via a site survey). These features could be  
15 used to either selectively filter-out content which falls outside of that criteria as the user  
16 navigates the information or identify if/when pages encountered where these user rating  
17 features are present, thus displaying this user statistical information in conjunction with the  
18 ratings statistics and/or associated annotations if desired.

19 We can allow the user to use one or more organizations to provide features and annotations  
20 to help in a search and information interpretation process, to add appropriate filter and bias  
21 to information presentation as a user browses the Web. Similarly, a user might adopt the  
22 endorsements of a friend for annotations, and editorial content about particular pages.

23  
24 The availability of feature information about individuals that provide high ratings for  
25 particular Web pages can also be useful for vendors that wish to find appropriate  
26 locations for their advertisements. For example, given an "ideal profile" of a consumer  
27 for a particular product, a web page/type of information can be determined that will be  
28 useful to find appropriate customers. This is an alternative method to providing adverts to  
29 particular individuals, instead choosing to provide adverts to limited information domains  
30 that attract appropriate users. Vendors may request additional information, such as site-  
31 specific page view correlations (including time spent viewing each page) in accordance

1 with user features. Exemplar user profiles and attributes of those users are certainly of  
2 interest to vendors to which those sites belong as well as affiliate sites on which their  
3 advertisements and/or syndicated products are advertised and sold remotely.

4  
5 We can also profile users with context information, for example based on their assumed  
6 goals (social, business, personal, etc.), information which can be determined by the nature of  
7 their current activities.

8 In an interesting extension we can allow users to contact other users that are browsing  
9 similar material with similar profiles in real time, allowing ad-hoc virtual work group  
10 formation. Of course, all of this is done within the carefully managed profile and privacy  
11 managed systems of SDI. The location of a user could play into decisions about physical  
12 meetings. Typically some users may wish to enter into a present (or future) dialogue, which  
13 can be scheduled via calendar agents, which perform automatically scheduled compatibility  
14 meeting/introduction functions.

15 It is also reasonable to provide the technique of collaborative filtering to identify users  
16 whose profiles are particularly similar to the material which is being viewed, e.g., those  
17 users who are determined to have a high proficiency level with regards to that particular  
18 material. Because the requesting user may be seeking to query the expert or seek edification  
19 on that material, typically same consideration is conveyed in exchange for that information.

20 In a variation, a more ----- data exchange may be provided if collaborative filtering  
21 identifies fairly similar users to the presently viewed material which also demonstrate, via  
22 their profiles, considerable complementarity in the particular knowledge which they are  
23 likely to possess. Vendors may also ----- and facilitate this type of  
24 information exchange on their sites between these complementary, (yet metrically similar)  
25 customers.

26  
27 In a consumer-vendor commerce application of this form of interactive browsing we can  
28 allow users to enter into dialogues with vendors based on their profiles, and share  
29 experiences with vendors, regarding their experiences. This can be done in real-time, or  
30 statically within an ianworthit style system architecture, where a user receives financial  
31 incentives for providing information that is useful to vendors.



1 what other clusters. Because conceptual flow in human language is quite complex and  
2 perhaps somewhat relative to the individual, it is useful to provide different interface  
3 settings, e.g., for example one which is tailored to the individual (i.e., the layout of the  
4 conceptual flow based upon statistics taken from all individuals which are generally similar  
5 to the interest profile of the present user what it would be for certain types of user profiles as  
6 explicitly indicated or what it would be for the collection of user profiles of those  
7 individuals within the present discussion groups (or if two individuals, the user profiles).

8 The large statistical information stored in SDI may be useful for the purpose of guiding  
9 discussions in a chat or forum context. Individuals may use the system in order to guide the  
10 flow of their own thoughts (where it is tuned to their own personal profiles) or that of what a  
11 particular individual with which they are corresponding at that moment (or in this case  
12 ideally both of their profiles) or (to appeal to the overall discussion forum or chat room) the  
13 collection of user profiles in that discussion forum. Again, in a variation, the system may be  
14 tailored to give additional weight to the particular historical experiences of the individual or  
15 group, thus if implemented in this way to the individual, the system may act in a similar  
16 capacity to a remembrance agent.

### 18 **1.3.11 Example: Vacation Package System**

19  
20 A vacation package organizer decides to begin a large-scale marketing campaign to target  
21 those people who would be the most interested in joining a new Caribbean Cruise. Although  
22 the vendor has a database of current customers, it is interested both in increasing the number  
23 and suitability of its potential leads.

24 Interfacing with the secure data interchange with which it is a member, the organizer  
25 identifies several possible sources of supplemental data: a LEIA-based travel discussion  
26 group, an on-line bookstore, and a Caribbean restaurant. These are found both by browsing  
27 through the interchange's internal list of members, and by using SDI-based data analysis  
28 tools, used within the interchange to automatically identify entities sharing common  
29 characteristics.

30 The package organizer then contacts each of these entities through the interchange, and  
31 negotiates different data-sharing deals: the travel discussion group is willing to exchange









1 by defendant , it such a equal mixture of empathizers from both sides constitute a  
2 virtual jury.

### 3 **1.3.14 Example: Group Therapy**

4  
5 There is a very useful and appropriate application for SDI to the automatic matching of  
6 individuals for purposes of group therapy. Individuals with commonality across many  
7 criteria can be selected, and multiple long-term groups adopted. In an application of this  
8 system, it may be possible to broadly extend the present scheme to the general public (where  
9 privacy is secured via the proxy server) and individuals may identify a group(s) which best  
10 fits their own unique emotional needs. It may also be useful to archive the sessions, index  
11 enabling the sessions and segments thereof to be searchable by keyword, e.g., via speech to  
12 text techniques and/or browsable by topical segment (which may be automatically  
13 segmented and labeled). The present system may also provide an ideal framework by which  
14 psychologists may identify patients of other psychologists that similar pseudonymous  
15 complements of psychological characteristics and symptoms, to allow targeted clinician  
16 interactions and robust and useful information from therapeutic approaches and/or drug  
17 treatment regimens.

18 In a variation, the present system methodology also enables a means by which much more  
19 specialized group session topics may be created which focus upon a very specific type of  
20 disorder, conflict or aspect of the patient's psyche. In this way, it may be possible for  
21 "identical" patients to, as group, focus upon different aspects of their ideal therapeutic  
22 regimen which collectively create a comprehensive customized treatment program. It is  
23 worth noting that the specialization possibly associated with particular focused sessions (and  
24 even types of individuals) provides a framework by which clinicians can become extremely  
25 specialized and expert within certain specific sub-domains of the field.

26  
27 It is certainly possible to adapt the above described technique for "smart interface" to  
28 provide useful ideas and suggestions for patients engaging in either psychotherapy or group  
29 therapy thus leveraging the information and personal experiences within the therapeutic  
30 processes of the collection of patients precoding them within similar topical sessions which  
31 shared very similar psychological and pathological patent profiles.

It is certainly reasonable and appropriate to adopt the present application framework to more several fields of medicine. For example, enabling physicians to identify other physicians have ----- had patients with the identical pathological profile and medical history. This enables them to ----- current medical history and clinical data, insights, observations, etc., in order for the physician to ----- the present clinical based on the expert advice and collaborative feedback from the other physicians possessing very similar experiences. Likewise patients could instantly access the physician or physicians which have or had experienced the most similar clinical situations that and the present situation of the patient in order to seek a second opinion. Physicians sharing unique clinical experiences (particularly clinically or scientifically in----- or series/problematic could use the ----- ----- scheme in order to ----- ----- companies, etc.. Finally one could also apply the -----  
-----  
-----.

### 1.3.15 Example: A Personalized Educational Portal

Another interesting application is to develop and deploy a personalized on-line informational portal containing everything from helpful links, illustrative content, text book information, quiz questions etc. The Personalized Educational Portal (PEP) typically is designed specifically for a particular class and includes several key features which are optimally applied as part of a comprehensive intelligent educational system. These features include the following:

- The ability to dynamically customize all forms of relevant information from the educational portal. The criteria for this customization however is based not upon the preferences of the user (unless for example the user receives credit for studies or projects or research on topics which s/he may select, rather it is based upon a predicted profile of the user reflecting his/her strengths and in understanding the relevant content. In particular the techniques of the issued patent number 6,029,195 [WHAT IS THIS PATENT NAME] describes a variation of user profiling in which

1 users are able to achieve a proficiency profile within certain domains of  
2 informational content where these informational domain(s) are determined in  
3 accordance with user's ability to answer a certain question(s) intelligently, discuss  
4 the answer to a certain question or about a certain topic or provide a useful reference  
5 or URL based upon the level of satisfaction of the requestor. This technique in itself  
6 could be usefully applied within the present application framework. For example,  
7 users may be students, and "experts" fielding questions could be other students (like  
8 tutors of sorts) and the payment they receive may be monetary compensation or even  
9 school, credit where other student's satisfaction ratings both qualify them for future  
10 opportunities to submit future responses within that particular knowledge domain  
11 as well as means of verification, and measurements of his/her proficiency over that  
12 particular material.

- 13 • Using the techniques of the parent patent application it is possible to also customize  
14 content delivery (including quantity, depth, and difficulty level) which reflects the  
15 user's strengths /weaknesses within the various relevant knowledge domains. It is  
16 possible to even statistically correlate the user's correct /incorrect responses to  
17 certain questions or types thereof as correlated with each other by common terms in  
18 the question answer part concurring in similar textual segments. These questions  
19 may be found in on-line exams or quiz questions associated with the actual content).  
20 It is possible to statistically correlate which content tends to best remedy certain  
21 deficiencies (as determined by incorrect responses to certain questions) by the user  
22 importance in being able to readily correctly answer those types of questions  
23 following reading the pages.

24 It may also be useful to apply the above technique towards determining which exam  
25 questions or combination thereof most commonly are revealing (if the student answers them  
26 correctly) of a student's proficiency within a particular small yet well definable sub-domain  
27 of knowledge. i.e. in which statistics show that If the student answers that sub-group  
28 correctly, s/he will likely answer the others in that sub domain correctly as well.

29 A sub-set of each of these exemplary questions from each knowledge sub-domain may be  
30 allocated as questions provided at the end of each relevant section of reading content  
31 (which may be on-line, off-line or only the responses may be entered on-line, exclusively

1 and if a wrong response is submitted, the system may recommend reading certain content  
2 which has (as above suggested) been statistically demonstrated to improve the student's  
3 proficiency in that sub-domain. It may also be possible to gauge in advance the students'  
4 predicted level and rate of advancement by knowledge domain via on-line psychological  
5 testing. Statistical correlations between these test questions and the student's learning  
6 abilities (by knowledge domain) can be usefully developed.

7 The system could even be constructed hierarchically by graduation of skill levels i.e. initial  
8 mix of questions of varying difficulty levels in order to gauge what level of content to  
9 recommend. The questions which follow are one level higher as is the next set of  
10 recommended content unless improvement is not observed.

11 It may be possible to a decision tree which automatically select which questions most  
12 effectively test the students, command of the materials (in this way a shorter exam may  
13 actually be better than for longer exams those questions which are more exemplary of such  
14 knowledge may carry greater weight). We could also effectively try to create a summary  
15 e.g. by topics /headings or even using text summarization techniques submit the areas of  
16 deficiency in one student to another student who is particularly proficient in that area for  
17 virtual tutorial. For this the tutor can receive monetary and/or also scholastic credit. The  
18 ideal scenario is actually creating virtual study sessions in which students are matched  
19 together which have the most different complement of proficiencies/deficiencies as possible  
20 and where these areas of potential mutual exchange benefit are maximized and are revealed  
21 to both students initially. Other forms of virtual study groups could be achieved by  
22 grouping students by similar proficiency areas and levels and creating a "virtual tour" with  
23 questions and content presented synchronously to a group of correspondence enabled  
24 students. In this way, if virtual tutoring or study groups occurs on-line exclusively, it may  
25 be performed pseudonymously as desired. We can even take a defined study group and  
26 assign it to group projects which can (again) relate to content to which the group is  
27 (collectively) deficient. (as group oriented problem solving has been proven to be extremely  
28 effective in remedying such deficiencies.

29 In a very novel application, it may be possible to even statistically using historical data to  
30 analyze the relative proficiency profiles of students by knowledge domain (perhaps) even  
31 further by, learning ability in those domains by teacher. Thus a virtual class could be

1 constructed using live streaming video which are truly customized to each student's  
2 particular needs and abilities, i.e. the idea would be before the semester to identify which  
3 teacher(s) were most successful in eliciting the highest levels of achievement among  
4 students which have an identical profile to that of the student (such statistics used in this  
5 profile may include but are not limited to the proficiencies / grades of the student in all other  
6 previous classes, psychological testing and/or (selective) knowledge domain proficiency  
7 questions (as above described).

8

9 Certainly depending upon such investor's personal preference, it is reasonable for each to  
10 agree within their personal data disclosure policies to disclose to SDI data regarding the  
11 types of investments, amounts and under what conditions (including those conditions stated  
12 by other investors such as the example above). The conditions for this disclosure, however,  
13 would be that no individual data be disclosed to another investor, i.e., that it be used only for  
14 SDI to be able to reveal aggregate statistics and predictions how the investment community,  
15 in aggregate, interacts with investment opportunities, and under what conditions, and in  
16 exchange for this disclosure, each investor earns the right to access this aggregate data. The  
17 individual investor or SDI, acting on behalf of that investor, may then suggest the best terms  
18 and conditions with which to negotiate with other investors, in order to optimally achieve  
19 their own desired objectives, part of the condition statement to SDI could also include the  
20 amount of money to invest in each synergistic investment opportunity which SDI typically  
21 would suggest to the investor based upon the relative distribution of its presently invested  
22 funds, size of the available non-invested funds as well as the relative risk/benefit proposition  
23 of the company to that investor (as estimated by SDI's use of company data as disclosed to  
24 that investor). SDI can even recommend or act as a proxy, in negotiating with the individual  
25 companies (who themselves could use an SDI negotiating agent representing their own  
26 interests). E.g., SDI, acting for the investor, could leverage considerable information about  
27 not only, the present investment opportunity, but literally all other investment opportunities  
28 which had been submitted to SDI for consideration by SDI affiliated investors. Based upon  
29 the total benefit which the investor stands to gain which is based, in turn, upon the SDI  
30 recommended comprehensive investment strategy) and the relative downside (relative risks





accept from the collective group of investors. Both sets of information (typically recommended by the SDI representation each associated party) are disclosed to the main (party neutral) SDI entity. If overlap exists in the range of terms or the ranges are close to one another, the introduction and negotiation between the parties becomes initiated (through the company(s) are not made aware of whether the overlap exists or is merely "within the general proximity" of accept-----  
-----introductions on the interest of the parties (though most markedly the companies)

### 1.3.16 Example: A "Group-Think" Ideas Market

In the consumer driven market place there is a significant unrecognized opportunity in being able to harness the intellectual capital of the consumer market. Historically, (in a non-networked environment) leveraging of such knowledge has been impractical and virtually infeasible due to the inability to identify and measure this resource of human skills (and thus match this skill with in the appropriate matching problem)with the particular types of problems which companies may be in need of new ideas and solutions as well as the lack of a trusted intermediary which can validate on an impartial including basis the authentic contributions of the provider of the information (as not having been claimed by the recipient after the fact as having been previously conceived before). Of course, there are remaining obstacles such as the fact that important problems are sometimes maintained confidential within a company (not even disclosed to some its own internal staff).

With the emerging of extranets a certain amount of inter-organizational knowledge leveraging is becoming more readily achievable. A secondary advantage over an extranet is because of its ability to act as a trusted intermediary. It can validate the fact that that the individual is only dealing with one commercial entity and not its competitors (at least formally). Thus unlike extranet based knowledge leveraging it is conceivable a trust relationship enabling a certain amount of privileged information disclosure is at least theoretically possible.



(as similarly) may enable the automated or semi-automated selection of disclosed segments of users (containing particular attributes) which are likely to have relevant properties. The user selection process may be iterative based upon responses from users as they are selected. We can demonstrate this value experimentally, for example we can offer a vendor a free-trial and present personalized information/advertisements to one group of SDI users (both on the site and if desired also across the network), and regular advertisements etc. to another group without the aid of SDI. The increase in vendor revenue can be estimated from client-level monitoring of the change in purchase volume achieved with well-focused solicitations **on the vendor's own business**. Other prospective ads and syndicated purchasables could likewise be virtually overlaid on the vendor's site. Thus degree of increased click through rates, transaction rates and syndication revenues could be tabulated for that vendor in advance of his/her subscribing (or even being solicited) to SDI. Portals are also potential beneficiaries of this service in light of their need to better target their partner's ads. The ability to the tailor a targeted marketing strategy to users knowing these correlations could be extremely beneficial in both on-line as well as off-line commercial environments. In the process of selecting the target user profiles significant value may be achieved from data collected from that vendor's own on-line presence, and a portion can be sold to other similar vendors; as discussed in the main description of SDI, this now forms a cached query which can be used by other vendors.

We might also demonstrate value to vendors with SDI by monitoring the performance of vendors with SDI technology, and providing metrics for new vendors to allow them to select suitable models of user targeting and personalization. Vendors that subscribe to SDI (the vendor centric version) can provide more attractive offers/products to users, based on information about the wider activities/interests of a user, on other vendor pages, and in the physical world (of course, only to the extent that this information is authorized by the user). Vendors can use information in the central SDI data warehouse for users' cross vendor and within vendor browsing and purchasing habits, and also with respect to profiling information about a user.

*One key application is first-time personalization*, so that information and products can be targeted to a user when he/she first hits a web page, based on profile information that the

1 user is willing to release. Vendors and users may also sell this information to other  
2 vendors. For example, information that a user likes a particular type of music is very  
3 valuable for vendors that sell content-based products, for example books, and CDs.  
4 Vendors can personalize their service, for example with collaborative-filtering based  
5 recommender systems.

### 6 **1.3.18 Example: An Efficient Product Delivery System**

7  
8 The parent issued patent [FILL IN THE PATENT NO.] describes an application of  
9 collaborative filtering to the *strategic optimization of a vendor's business*, for example to  
10 allow a vendor to select an optimal location for an inventory warehouse based on projected  
11 consumer purchasing patterns; using the aggregate purchase history of users at that site  
12 compared with the other purchase selections at other sites for similar users. The model can  
13 also be used to predict demand for new items, and optimal locations for inventory given  
14 warehouse locations.

15 Now, within SDI the vendor can have control over two things:

- 16 1) the vendor has access to profile information about users
- 17 2) the vendor has a method to personalize and annotate on a dynamic basis the  
18 presentation of products to a user.

19 In common with the *smart caching* application of SDI to making communication networks  
20 more efficient, we can suggest an application to a product delivery system. At any moment a  
21 seller can have a good idea of the products that are available, and even the products that are  
22 in trucks etc., close to a prospective customer. This information allows the *vendor to*  
23 *selectively present items that are close to the user, with labels such as "this is within*  
24 *30minutes of your door, NOW"*. Messages like this allow a vendor to manage its delivery  
25 process, because predictions can be made about likely purchase patterns, and then stock that  
26 is brought close to user locations can be explicitly advertised to those users.

27 Essentially, we attempt to predict transaction volume on a per item basis, and then  
28 positioning geographically physical inventory storage facility locations where inventory can  
29 be stocked so that items which a user is predicted to request are already located within the  
30 immediate physical proximity of that user.

1 As the trend to disintermediation continues (whereby we have informational middlemen, but  
2 not middlemen that physically hold products) we can allow vendors to notify portal  
3 intermediaries of product location within SDI, so that the portals can then pass this  
4 information onto customers.

5 The key role of the SDI system in this application is to allow users to receive personalized  
6 recommendations of products, and also product locations, based on their profile  
7 information—but without the vendors receiving that profile information. In addition, the  
8 shared information about user profiles in the data warehouse allows vendors to build good  
9 models for inventory location. Finally, SDI can act as a trusted intermediary between  
10 different competing vendors: for example suppose Amazon users bookstore A and  
11 bookstore B to provide books, and A and B do not wish to tell each other where their  
12 products are. The stores tell SDI where the books are, and SDI presents that information to  
13 users on a per-book basis, as books are browsed within the purchasing system.

14 There are numerous useful applications to improving quality, speed and cost of delivery to a  
15 user. In one example, it may be possible to provide same day delivery for on-line purchases.  
16 We could also restock a truck on the basis of what a local population of users are likely to  
17 purchase; with items presented to a user along with a particular anticipated delivery period,  
18 with the Global Positioning System (GPS) or more LEIA [INSERT THE PATENT NO.]  
19 on the vehicle providing up-to-date information about a truck's location and anticipated  
20 delivery time.

21 We might also suggest that a user can physically travel to some location close to his/her  
22 base, based on information within SDI about the user's travel patterns, etc. Purchasable  
23 items can then be located in locations that are convenient for users, or shipped dynamically  
24 to those locations (i.e. intermediate warehousing systems). Users can be provided with real-  
25 time directions to the location of such a pick-up point via personal digital assistants (PDA)  
26 and navigational devices, etc.

#### 27 1.4 Push-based Advertising/Solicitation

28 Another form of information that a user can submit to the central SDI database via his/her  
29 client is the advertising acceptance functions, which state the explicit preferences of a user  
30 for adverts, and required payments. The level of compensation that a user requires to receive  
31 an advertisement will depend on the relevance of an advert.

1 An application of SDI is that vendors can use the centrally stored information to request that  
2 as a result of a query the query-execution module in the central SDI data warehouse sends  
3 advertising solicitations to appropriate users, i.e. users that are willing to receive adverts as  
4 specified within their acceptance functions, and willing to release profile information to  
5 enable useful advertising as specified in their price rules for regular data.

6 Figure 15 illustrates the system for push-based advertising. A vendor requests the right to  
7 advertise to users by sending a special type of query to the central SDI data base, where the  
8 query states the vendors preferences for user profiles and requests that adverts be sent to  
9 users. The vendor never receives the contact information for user's directly. The server also  
10 receives payments from vendors and credits users' accounts.

11 Suppose Vendor V wants to advertise a new product to relevant users. Vendor V can use  
12 the SDI database to select users with a high predicted hit rate for the advert, and then  
13 compute the value of a bid that it is prepared to pay users for the right to provide them with  
14 its advert. The vendor makes a request for bids to appropriate users by formulating this as a  
15 query with an associated action, and submits to the central SDI data warehouse where it is  
16 handled by the query-execution module. Part of the query is a bid function, indicating the  
17 maximum amount that a vendor will pay to display an advert.

18 The query-execution module evaluates the request with respect to the profiles of each user,  
19 the rules that agents submit for allowing access to profile information, and the advert  
20 acceptance functions. If successful the vendor pays the minimum value necessary to have its  
21 ad accepted by the user, again simulating a Vickrey auction as for the dynamic competition  
22 for banner ads described earlier.

23 For example, suppose a vendor has an advert **Ad** to push to users, and is willing to pay up to  
24 \$1 to place the advert with users that have a predicted hit rate of greater than 30% on the  
25 advert. The vendor can formulate a query that will first identify users with a predicted hit  
26 rate that is greater than 30%, and then determine which users will accept the adverts at the  
27 bid price. If successful, the adverts are pushed to the users via the SDI central data  
28 warehouse, and the payment required by each user is submitted, less than \$1 in all cases. In  
29 a simple variation, the vendor might also state a fixed budget, so that it does not send  
30 advertisements to more users than it can afford.

31





- 1 (i) The vendor assesses the value of the information present in the secure data  
2 interchange. This computation is performed securely either by revealing randomized  
3 aggregates to the vendor to enable its own local analysis, or by allowing the vendor to  
4 check data and algorithms into the secure data interchange site for analysis.
- 5 (ii) The vendor selects criteria for mailing unsolicited advertisements, and agrees on a  
6 pricing model. In this case per-impression pricing is the most obvious pricing model,  
7 as it is difficult to monitor when a user responds to unsolicited mail per-transaction  
8 pricing is difficult. The user could be motivated to do this should the Secure Data  
9 Interchange promise future returns for recording a successful solicitation with the  
10 database.
- 11 (iii) Either the data list is released to the vendor for its use, if this is within the selling  
12 vendor's data policy, or the data interchange sends mailings on behalf of the  
13 purchasing vendor.

## 14 1.5 Community Dollars and Business Models

15  
16 The primary objective of the iamworthit model is to create a market for information  
17 about buyers, with agents that submit profile and other personal information to the  
18 database able to collect payment in return for queries executed by vendors. In an  
19 important extension, we can allow this payment to be made in terms of *community*  
20 *dollars*, which can only be spent with particular vendors. Community dollars allow a user  
21 that provides profile information to the system of SDI to receive payments that are  
22 *dedicated for a particular type of purchase*, these payments are called "community  
23 dollars". One central example is to allow a vendor that signs a user into the SDI system to  
24 be able to lock a proportion of payments accrued by that user to the vendor's own  
25 product/service domain. This is important, because vendors now have incentives to bring  
26 new users into the system. A vendor can offer a user community dollars on its (and its  
27 affiliates) web site, in exchange for receiving user data via the SDI data exchange.

28  
29 We allow vendors to pay in "community dollars" for adverts, dollars that can only be spent  
30 at that vendor (with the host site of the advert receiving a share of the profits). This provides  
31 vendors with the ability to gain long-term customers. Furthermore, so long as the buyer

1 agrees to receive advertising from his/her iamworthit subscription offer, community dollars  
2 can be replenished at the rate at which advertisers are willing to pay for impressions. This  
3 provides buyers with an incentive to spend at the vendor's site, because the vendor can  
4 monitor (pseudonymously) the buyer's that are sensitive to discounts and other special  
5 offers (that are delivered as community dollars). With community dollars a vendor can  
6 compensate buyers for information that they access, but tie that information to certain  
7 vendors (e.g. the vendors that first signed the user into the system). The system has the  
8 following useful properties:

9 (1) buyers are incentivized to provide information that allows vendors to push relevant  
10 advertisements/products;

11 (2) buyers will also be more likely to make purchases at a site or a coalition of sites for  
12 which they can receive discounts via community dollars;

13 (3) providing buyers with community dollars will increase the number of hits to a site.

14 We allow buyers to receive compensation for providing personal data to vendors,  
15 information that has value to vendors because it allows information to be focused (for  
16 example relevant ads can be displayed to a buyer, based on his/her profile). The system of  
17 iamworthit credits buyers for information, and provides buyers with direct incentives to  
18 reveal profile information to vendors.

19 A vendor can sign up with iamworthit.com and agree to provide only the most restrictive  
20 type of community dollars, that can be spent at that vendors site. Community dollars are the  
21 currency that vendors provide in return for the right to provide focused information to  
22 buyers. Dollars can be general (e.g. for a network of vendors), or very tightly focused (e.g.  
23 for a particular product, at a particular time). The buyer-centric infomediary acts as a broker,  
24 matching buyers and vendors. Another key role of the infomediary (e.g. the portal) is to  
25 protect the buyer from information saturation by controlling the flow of solicitations. (i.e.  
26 restrict the number of ads. that a buyer sees)

27 Community dollars extend the methods in Secure Data Interchange that allow a user agent  
28 to sell controlled access to information to other agents, introducing in addition a method  
29 which can control the ability of a user to spend received payments. For example, a vendor  
30 might be happy to pay a user for the right to display a targeted advertisement if the payment  
31 represents a discount that can only be used against products offered by that vendor.

1 Community dollars provide this functionality, allowing dollar payments to be limited to  
2 reimbursement in particular "communities" of vendors and/or other agents.

3 Community dollars instead are more configurable than traditional currencies, which by their  
4 very nature are transferable to any agent. In particular community dollars can be restricted to  
5 spending at certain vendors, can also have non-linear values to encourage loyalty from  
6 buyers.

7 The novelty in our description of community dollars in comparison to current methods  
8 known in the art for electronic cash is that we allow e-cash to be targeted so that it can be  
9 spent in particular ways, to be smart so that it can accrue value in non-linear ways, and to be  
10 transferable only between the pseudonyms that belong to the same agent.

### 11 **1.5.1 Simple Business Models**

12  
13 A vendor that allows community dollars to be spent does not need to implement a special  
14 community dollars/discounts program. The buyer can also be issued a special debit account  
15 dedicated to community dollars, that permits pseudonymous transactions without revealing  
16 a buyer's portfolio of pseudonyms.

17 A portal site that hosts advertisers and buyers that subscribe to iamworthit can mandate that  
18 all community dollars are to be spent at sites that advertise on the portal site, and also only  
19 when the sites are accessed via the portal site. This technique will increase portal traffic.  
20 Portals can be expected to compete in terms of: (a) the fraction of advertising revenue that is  
21 turned over to buyers, in return for receiving profile information from buyers; (b) the level  
22 of advertising that buyers are exposed to; © the nature of the community dollars "package",  
23 i.e. what vendors can the dollars be used at etc. This can be useful to attract niche customers,  
24 that have common outlooks, interests, and business needs. The primary goal of the portal is  
25 to drive traffic through the portal.

26 The main mode of the community dollars advertising model allows vendors to advertise for  
27 free, but provide community dollars to buyers, that can be spent at some later time. The cost  
28 of advertising can be linked to the success of advertising. Moreover, the vendor can direct  
29 offers and adverts to particular user profiles. The hosting web page receives a share of the  
30 vendor's revenue that comes from transactions involving community dollars. The dollars  
31 can represent "stored value", such as bonus points, that can be applied to special discounts

for offers which are delivered via digital coupons and/or as “straight value” which could be converted directly to purchases thus are equivalent to real dollars at the point of transaction. The community dollars can be “credits” that can be redeemed as real cash, credits towards discounts, and can be spent across a suite of sites, or limited to one site. The co-pending patent application entitled “System for the Automatic Determination of Customized Prices and Promotions” [INSERT THE PATENT NO.] describes a comprehensive scheme which may be implemented in either on-line or off-line commerce environments. The system enables vendors to deliver a digital message in the form of a promise to a buyer (typically on encrypted form for purposes of targeting a buyer specifically). This promise is typically a discount for a product, set of products (or all products in stock) or may even include entitlement to special privileges for that buyer, thus it is termed a “digital coupon”. The community dollars can represent special discounts for a buyer.

The buyer receives a financial incentive for receiving well-targeted solicitations, while preserving buyer privacy within the SDI system. The vendors support the community dollars through advertising revenues and increased sales volume. We can also provide the vendor through which the buyer first subscribes a special “first screen” right that allows the vendor to provide a buyer with his/her first impression as soon as s/he logs on.

Of course many variations of the community dollars scheme are conceivable and the current description is in no way intended to limit the scope of the claimed invention. For example, the ad revenues generated may instead be apportioned between direct payment which the buyer receives, in real cash, community dollars which the vendor (or ISP) credits the buyer and/or direct payment which the vendor (or ISP) receives as well as any worth it or any combination of the above such as exclusively direct payment which is apportioned between the buyer and the vendor (in the absence of the community dollars scheme).

In one variation all community dollars collected by a buyer must be spent back at the vendor site at which they originally subscribed (and also the site that hosts the adverts of other vendors). A buyer can spend the dollars with any vendors that are site partners of the original site. This provides the vendor an incentive to accept and promote the community dollars concept.

The value of providing a buyer with targeted solicitations is estimated at approximately \$300 to \$500 per year (based upon \$120 per 1000 targeted impressions at approximately 25





which is possible via the one-way identity look-up functionality of the central SDI data warehouse. The eBank can query the look-up table, and check that both PIDs correspond to the same UUID.

For example, an amount of dollars below a certain threshold can be programmed to have no value, but additional dollars can have increasing marginal value. This provides a non-linear incentive for a user to stay with the same vendor in a market with low switching costs, or a non-linear incentive for a user to receive adverts from the same advertising agent. The advantages to vendors and advertisers are repeat purchases, consumer lock-in, and also improved profiling from interacting with the same agent over extended periods of time.

The functionality is embedded into the community dollars. The framework expands the idea presented in Chaum [Chaum 85; Chaum 92] where the public key of a public key /private key pair indicates the dollar value of a community dollar. In this case, we allow any number of public-key/private-key pairs, and use the public key as an index into a table maintained within the eBank to provide additional criteria about how the cash can be redeemed, for example it has more value if used in combination with other community dollars, it decays over time from date YY/YY/YY, etc... We do not intend to limit the variations possible, but provide this as a framework for ways in which constraints and conditions on community dollars can be extended. For example, community dollars can also be programmed to lose value over time. This can increase the number of repeat purchases

In the preferred implementation we use an “electronic cash” infrastructure for the community dollar system. A buyer’s SDI-enabled client-level proxy stores dollars that the buyer receives securely. Dollars are anonymous and non-traceable, so that the buyer can maintain a single “bank” of dollars, and aggregate dollars collected across pseudonyms for a single purchase, so long as the purchase satisfies the constraints on the dollars. Each dollar is created using Chaum’s blinded signature technique, and also signed with the conditions on its use.

This scheme allows vendors to monitor the offers that buyers respond to, because when a buyer presents a community dollar—the dollar can be validated to indicate the type of discount that it is, even if the identity of the dollar (i.e. the serial number) is untraceable. SDI provides vendors with guarantees that buyers have once-in-a-lifetime pseudonyms, so redeeming a voucher of a particular type that is redeemable only at vendor  $V$  and was issued







1 transactions. There is considerable economic pressure on web tool providers to provide  
2 competitive solutions which are not only rapidly deployable but also extremely robust  
3 (typically bundling some form of personalization technology), thus as a result of  
4 economic pressure to drive advertising and e-commerce transactions through the site  
5 (often through the use of a revenue sharing model which can reduce upfront costs to the  
6 vendor substantially). iamworthit through its customer loyalty enhancement capability  
7 provided by community dollars can significantly improve the effectiveness of  
8 personalization at the site (perhaps more significantly than cookies which often are  
9 blocked and all typically not implemented so as to be recognizable across the collection  
10 of sites which the tool provider supplies or the web host services).

11 It is believed that the revenue sharing model may be structured to cover both the cost of  
12 the tools, web development services as well as web hosting (thus unless he himself  
13 offers a complete turn key suite of solutions and services a tool provider which  
14 integrated the community dollars concept could be in an extremely competitively  
15 advantageous position in attracting strategic relationships with Web developers and Web  
16 hosts. The tool provider may for example sell an ad server module or could  
17 "automatically subscribe" the user with an opt-out option. The very fact that all of the  
18 tool , Web development or hosting provider's sites are exempt from ad blocking/ad  
19 replacement technology (barring iamworthit competitors) would be a significant  
20 motivating factor in incentivizing sites to utilize their services (e.g. through promotional  
21 programs). An emerging large market for Web development is providing these services  
22 to individual end-users. Completely free services in this regard would likely mobilize  
23 what activity in what is now a dormant, albeit enormous commercial market. For  
24 example, tools for developing elaborate Web-wide community portals could have  
25 customized community links which are matched to the user's personal profile. Visitors  
26 to the site (subscribed to iamworthit) could experience an additional layer of  
27 personalization based upon their own user profiles.

28 Web hosts may also use an additional feature (optionally to significantly drive increased  
29 traffic through potentially all sites on their network. This optional feature is a set of  
30 links (e.g. along the side of the user's screen). These links are used to point the user to  
31 other relevant pages issuing aggregate site usage statistics of their visitors (as described

1 in the parent patent application) iamworthit, user referral links may additionally  
2 personalized based upon their user profiles.

### 3 **1.5.4 Vendor coalitions**

4 Vendors may choose to form coalitions, to allow buyers to spend community dollars at  
5 any “partner” site. Vendors that have similar buyer bases can be automatically identified  
6 using collaborative filtering. (i.e. determining similarity with the present vendor, from the  
7 aggregate vendor preferences of a given vendor’s subscribers). Also, these resulting  
8 metrics could incorporate predicted online spending by each buyer at each site. This  
9 could help to narrow the selection of sites the vendor wishes to partner with and/or the  
10 selection of these partner sites could be determined and presented to the buyer to even  
11 further narrow the selection for each buyer. All vendors in a coalition advertise, and  
12 provide cross-links and up-links to other vendors.

13 The coalition model is good for buyers that are more likely to find products that they  
14 want. Vendors can share the risk of advertising, since dollars provided to one buyer by a  
15 particular vendor can be redeemed at another vendor. Advertising and community dollars  
16 increases sales volume at all vendors in the coalition. Furthermore, studies on on-line  
17 buying behavior suggest that on-line shoppers tend to make purchases across a variety of  
18 categories of e-commerce products, thus it is likely that they would also prefer the  
19 freedom and flexibility which is associated and can be provided in this way only by  
20 deploying a multi-site community dollars scheme. Iamworthit usage statistics are very  
21 effective in identifying prospective vendor coalitions which consist of complementary  
22 (non-competitive) vendors. Alternatively, personalized coalitions which can be created in  
23 ad-hoc fashion for each buyer can further serve the buyer’s interests by increasing the  
24 freedom and flexibility which the buyer often strongly desires. Though practically this  
25 can only occur for those (perhaps smaller less commercially significant) vendors which  
26 have not been a major coalition consisting of vendors which many buyers tend to use  
27 which collectively serve a substantial percentage of the overall user base.

28 Thus if the vendors have not established firm partnerships with other vendors, we can  
29 even allow vendors to form dynamic and virtual coalitions within SDI, with a potentially  
30 unique coalition of vendors for each buyer. The coalition may consist of an optimal pool

1 of vendors, as determined by SDI collaborative filtering techniques. The goal in this  
2 model is to provide buyers with a particular “brand” of community dollars.

3 • Multi-vendor Community Dollar Portal

4 We can allow each vendor to retain an exclusive right to advertise to each buyer; and  
5 also develop a portal for the coalition—that gives advertising prominence to coalition  
6 members. Portals will be expected to aggressively promote community dollars.

7 Buyers that collect community dollars become loyal return visitors to the portal and  
8 its associated vendors. In the case the vendors do not generate the same value we can  
9 provide community dollars in proportion to the value that a vendor contributes to a  
10 coalition.

11 We can also provide targeted advertisements for the vendors at the portal, using the  
12 user profile to focus ads. The categories and links at a portal (that might include a  
13 search engine) can be re-prioritized (highlighted and/or re-ranked) in accordance with  
14 the buyer’s preferences (as described above), and to favor subscribing vendors. In  
15 commercial practice, the motivation for vendors to become coalition members is  
16 largely driven by customer partnerships with providers of e-commerce tools and  
17 solutions. There are many emerging trends by which these intermediary tool  
18 providers could conceivably integrate community dollars. Consider loyalty points  
19 (e.g. [www.mypoints.com](http://www.mypoints.com)) or loyalty bonuses programs (e.g. [www.clickrewards.com](http://www.clickrewards.com)).

20 These points or bonuses could be substantially subsidized by the advertising and  
21 deeper benefits passed onto the consumer. Other simpler technology which would  
22 compel these vendors to cooperate include shopping basket technology, the emerging  
23 standard, ECML, common Web tool and/or tool development solutions, common  
24 hosting solutions, common ad delivery systems.

25 Vendors pay the portal site to advertise, and the portal provides community dollars to  
26 buyers in return for privacy-protected profile information. This model does not  
27 provide incentives for the portal to provide well-targeted adverts, because there is no  
28 direct link between a portal’s revenue stream and the vendors’ sales volumes.

29 A portal with community dollars that can only be spent under a single pseudonym at  
30 its partner sites also provides an incentive to buyers to interact under a single  
31 pseudonym—which in turn allows a portal to profile buyers across its complete



types of commercial partners) in order to dedicate a certain percentage of the community dollars (e.g. thirty percent or approximately one hundred fifty dollars per customer) which could only be redeemed at that multi-store retail site (and/or the value of these dollars could be worth more at the retail site). In addition, in this model, the independent advertising initiative of iamworthit would be geared towards community dollar credit of that retail site. It should be noted that, because if other outside competition occurs to the basic iamworthit scheme to a substantial degree there will not be a compelling incentive for buyers to adopt a more restricted form of valve (as retail credits at a particular site), versus accepting the credit from a competitor in the form of cash. Thus this model could provide a viable means for attaining a leading position in one or more on-line retail markets if this competition does not substantially exist.

- **Marketing Network.** Iamworthit sites which offer a community dollars promotion could, upon the buyers subscribing to iamworthit, additionally offer the buyer with a down-loadable client based software which provides a small promotion in conjunction with a link to iamworthit. Each time a recipient of the email subscribes to iamworthit, a percentage of the value of that customer is credited back to the buyer in the form of community dollars. Each subscriber resulting from the current subscriber's email (though reduced) provides an additional credit to the original subscriber in accordance with the marketing network business model. If the site originally delivering the promotion is not an e-commerce site, a percentage of advertising revenues resulting from the subscriber (and potentially all resulting subscribers) could be used. It could be applied in the form of iamworthit advertising (or exchanged) for advertising in an ad server.
- **Free Web hosting.** A portion of community dollars may be allocated in a revenue sharing arrangement between the buyer and the web-site from which subscriptions to the service are provided, thus enabling Web hosts to offer free hosting services while also receiving full payment which is iamworthit advertising supported. Iamworthit could also easily use its vast statistical data collected from each site in

the hosting network in order to predict which sites tend to be most commonly linked to from the present site (and/or share “similar content”). This technique is described in the parent patent application by establishing these links users can better access relevant information . In one preferred variation, all non- customer’s sites are not included in the referral links. This level of targeting is likely to drive significant traffic through the Web hosts’ network (and even more so if iamworthit user profiles are transferred to the site upon visitation of an iamworthit subscriber as these similar cross-links can also be personalized to the user.

### 1.5.5 Transaction-based Revenue-sharing

In transaction-based revenue sharing, the only time that advertisers pay to provide an impression to a user is when a sale results, in this case the hosting site receives a cut of the final transaction price. The vendors provide buyers with community dollars directly. The dollars, which are stored at the portal site, allow buyer-spending to be tracked. This allows the portal to monitor when a sale occurs, not just a hit on a banner ad. With transaction-based revenue of this kind, personalization is critical. In this model the portal will give prominence to adverts from successful sites. A portal site may forgo payment from a vendor in exchange for the increased click-through from a strong network of community-dollar enabled vendors. Value is credited directly to buyers for future redemption at that particular vendor's site. The community dollars provided to a buyer can be restricted, such that a buyer can only redeem dollars if s/he maintains enough visits to the portal site.

Vendors can offer discounts on their own products directly, instead of providing the portal with money for advertising. The vendor only pays to the extent that its advertisements are well-targeted. The vendor could also request special ad priority. A vendor that presents advertisements to a buyer offers the buyer discounted promotional offers for products offered by partner vendors, in exchange for subscribing to iamworthit and receiving targeted impressions. These offers are in lieu of community dollars, and can be provided by partner vendors—maybe in exchange for a right to a number of ad deliveries for the vendor’s own advertising purposes.





1 implement this monitoring, because no other system knows a buyer's portfolio of  
2 pseudonyms. The buyer can present its digital credential when visiting a vendor's site.

3 A vendor may wish to provide loyalty dollar credit; for example, it would be possible for  
4 vendors to offer buyer's credits if the buyer is a 100% loyal customer i.e. that she/he did not  
5 (over a specified period) do his/her purchases at the site of any competitor. For example,  
6 certain types of high value customers could be given considerable value in the form of  
7 credits or discounts as a result of demonstrated vendor loyalty. The credential can be time-  
8 stamped, to prove loyalty. It does not reveal any information about the buyer's other  
9 pseudonyms to a vendor, because many pseudonyms will exist that have not made any  
10 purchases from a competitor. Upon accessing the vendor's site, this credential may be  
11 presented to the vendor. One criteria for the above benefits could be that the buyer may visit  
12 a competitor site, and engage in interactions; however s/he should not transact with that  
13 vendor.

#### 14 1.6 Alternative Business Models

- 15  
16 • Offer discounted or free services in return for the right to access profile  
17 information.

18  
19 Examples: Free or discounted retail products with "niche" partners in each category;  
20 Free dial-up ISP (as an independent ISP or a service to jointly promote free access  
21 with ISPs); Free Cable and ISP service, Free pay-per-view (note that viewing patterns  
22 and the associated content could provide additional valuable user profile  
23 information); Free phone service (e.g. advertise subscription service on screen phones  
24 or audio ads from pay phones); Free prepaid calling card; Free print media  
25 subscriptions (magazines, newspapers); Free book clubs; Offer any combination of  
26 the above with "deep discounts" for each (this can involve \$350 community dollars  
27 per buyer or it may simply involve certain purchasing limitations per customer).  
28 Each vertical niche partner in exchange gets exclusivity within their own respective  
29 niches to target advertise to those buyers (e.g. retailers); Free access to sporting  
30 events; Free credit for casinos; Free lottery tickets; Free charity donations;  
31 Discounted hotel lodging; Monetary credit to a credit or debit card (either an



Alternatively, it would be possible to offer websites the ability to become Internet service providers where the interface to the ISP home page would essentially be heavily branded to that site or portal. Companies like GTE already offer a “Virtual ISP” service in which the content to the ISP home page is unique to the ISP while the network is provided by the virtual ISP service. This model would be particularly compelling for sites which are largely community oriented and have a potentially loyal customer base. Moreover, interestingly, many of these community sites are offering many of the services and capabilities that a full-blown ISP would offer from its home-page, e. g., a portal interface, links to high-quality content, chat/forums, e-commerce, commerce affiliate links, etc.

- **Free Community-based Content.**

Create premium content which is free to iamworthit subscribers, subsidized by revenue from profile information. Each iamworthit buyer would be granted free access privileges to the premium content on all sites which are part of the program.



A menu of different forums and chats can be displayed on each iamworthit member site. (the underlying methodology for which is described in co-pending patent application “Virtual Community Service for System for Customized Electronic Identification of Desirable Objects”). In accordance with this specification, a variation of the service involves the process for identifying individuals who most closely match a given category or target object. For example in the context of the present implementation a category or content, merchandise or a purchasable being specially promoted may be the focal point of a discussion forum or chat room, which is automatically organized by the Virtual Community agent. Accordingly, a portal (or in accordance with the present trend) a site with which a portal interface is integrated utilize the present techniques for generating virtual communities for each category or sub-category of content on the portal or for direct access into a forum or chat room which was automatically created around that particular site (as the target object used as the matching criterion). As described, the user may navigate a hierarchical menu of virtual communities which may be constructed automatically according to the methods described which involves communities assigned to category, sub-category,



Deliver through the back of sales receipts, kiosks or direct mail or on-line substantial purchase credit to retailers (e.g. grocers') customers, using the aforementioned technique of utilizing a unique URL to identify the vendor and/or promotion from which an iamworthit subscriber originally accessed the iamworthit subscription site (thus identifying for both buyer and vendor the appropriate denomination and/or terms of community dollars issued to the buyer). In the preferred implementation, a loyalty card is used to identify the buyer thus enabling the community dollars value to be provided to the customer at check-out as straight credit or possibly an enhancement to loyalty credit. The buyer may also be identified via credit card or alternatively a voucher (or coupon) could be printed from the buyer's computer or from a kiosk which is typically situated near the entrance to the store and which could be activated upon insertion of a loyalty card credit card (or associated authorization code) and could also be used to disclose the buyer's community credit balance. A unique identifier for that voucher or coupon is provided and non-tamperability measures are provided such that the buyer's community dollars account can be appropriately debited upon redemption. Preferably, a pre-determined value is specified on each voucher (which could be predetermined by the service or the buyer) or alternatively, the total community dollars balance could be specified on the voucher along with the buyer's name/address and redeemable only upon presentation of valid buyer ID.

- **Free ASP Services** – Web-centric applications are becoming an increasing central part of the e-business infrastructure. iambic could potentially enable organizations to gain free access to these Web-centric applications in exchange for iambic enabling their employees and customers (e.g. which may use some of these applications as well).
- **Credit to a User's Credit Card** – Many credit cards are tailoring promotions to enhance not only acquisition of market share but also loyalty of its users. The user-centric SDI provides an enhanced (e.g. rebates for expenditures) platform which could provide more efficacious loyalty enhancement marketing strategies

for card issuers. For example, a credit card user could be given certain credit towards the card (e.g. via a promotional offer) for becoming an iambworthit subscriber (which is largely ad revenue supported). An additional novel application could involve providing an offer for an additional benefit which could be provided if the user performs all of his/her on-line transactions using the present credit and (i.e. wherever the card is accepted).

An iamworthit pseudonymous credit card [e.g. LMP 94] could be provided whereby users collect credit for advertising (e. g., on billing statements for the ad) which the user receives as well as the advertising delivered via iamworthit's on-line (Web or targeted email) as well as pseudonymous physical mail, telemarketing calls using user profile pseudonyms and one all time or pseudonymous phone numbers. This user profile data consists in part of off-line data and on-line data.

- **Free Personal Portal for Individuals.**

It is reasonable to offer individuals completely free, Web design/development and hosting services which are offered and mass marketed. At the time this patent was written, Web developers were offering such free services in exchange for revenue sharing for advertising product syndication or e-commerce offering on the individual's site. Iamworthit could conceivably be deployed in conjunction with this commercial model whereby either a portion of the user's own iamworthit profiles are used to subsidize (or subsidize in part) the portal and the developer and/or revenues shared from profits from iamworthit subscribers who subscribed from the portal are utilized also (or instead). As such, it is also reasonable for such a developer to provide links on affiliate sites (e.g. access an affiliate network) which offers free individual portals in exchange for subscribing to iamworthit (and perhaps agreeing to offer a similar iamworthit expectation from that portal as well from which the user can also share in a portion of the revenue stream. Such a program could even be structured in its revenue distribution to the subscribers as a multi-level marketing network. Users could even receive value for providing links to the vendors site either







1  
2 The mechanism for providing value to the recipient (in exchange for providing the  
3 desired action) is flexible and may include (as suggested above) “upside” value, direct  
4 compensation (or direct “redeemable value”), or an exchange of actions between agents  
5 (I will do X if you will do Y). Part of a deal may be that the provider of a requested  
6 action guarantees exclusivity and that it will not provide a similar action for competitors.

7  
8 A vendor might identify and propose deals to the customers of another vendor, to apply  
9 economic pressure through threatened boycott and compel the vendor to perform a  
10 desired action. This “indirect” economic leverage may be applied at the consumer  
11 (lowest) level in the supply chain, and passed upward. In addition to boycott, a vendor  
12 could promise please note, shouldn’t consumers be termed customers of the vendor since  
13 consumers cannot be threatened with boycott) consumers credit in return for taking a  
14 particular action. SDI can mediate the entire electronic control, transfer, fulfillment and  
15 negotiation processes (typically XML-based) individually on behalf of each entity within  
16 this multi-step negotiation process. In the above applications SDI can represent each  
17 entity in the negotiation process, and also have a role in providing information and  
18 statistics, expert opinions, etc. which clearly elucidate the value and potential value  
19 opportunities available to the entity(s).

20  
21 The current model allows different supply chains to compete for the loyalty of the  
22 consumer, e.g., through providing indirect economic value to the consumer as in kind  
23 products/services and/or upside in the benefiting entity.

24  
25 Economic alliances can be created between vendors, suppliers and consumers within the  
26 same supply chain. SDI can act on behalf of different supply chains and establish a  
27 competing market between different chains for the business and ultimate loyalty of  
28 purchasers. End purchasers have a large economic leverage on a supply-chain. The  
29 decision of an end purchaser will depend on:

- 30  
31 1) The type, selection and quality of personalized products/services



1 With interwoven supply-chains, for example, with the same supplier connected to  
2 multiple vendors, and vendors with competing suppliers, this method of passing value  
3 along the supply chain becomes more tricky. When a consumer applies economic  
4 pressure on a vendor that it deals with in the supply chain, that vendor may inform other  
5 suppliers of the pressure and block service to the user. A bartering system can provide  
6 additional benefits, to allow vendors to exchange bonuses, discounts, etc. with favors  
7 offered to vendors in other value chains. In general it is quite likely that a favor will not  
8 substantially benefit a user in another supply chain. We can also include a method to  
9 prevent a vendor that does not itself offer favors in the supply chain from taking  
10 advantage of favors offered by other vendors. This can support cooperation  
11 within the supply chain, and allow effective competition with other supply chains.

12  
13 Each vendor in the supply chain can benefit in two ways:

- 14  
15 1) They receive a “commission” on the transactions in which they convey the favor  
16 to their customer (which again may be “upside” in the ultimate beneficiary of the  
17 favor and/or product/service loyalty credit with their own supplier).
- 18  
19 2) (Most importantly), they receive customer loyalty advantages via the community  
20 credit they provide such customers.

21  
22 For reasons of this latter benefit they are likely to compete with other competitive  
23 vendors for the privilege of exclusively delivering the incentive to evoke the favor to the  
24 next level up the supply chain. In this way each vendor below the supplier requesting  
25 stands to gain via customer loyalty benefits by competing with another supply chain.

26  
27 A consumer can provide value to a supply chain by exclusive provision of personal data,  
28 to allow a supply chain to create specially customized and targeted products and services.  
29 However, it is not necessary for a customer to have an exclusive contract with a supply  
30 chain (although this might bring greater loyalty benefits and payments).

1 A variety of different techniques and methods should be integrated and deployed to  
2 implement this agent-mediated supply chain value exchange system:

- 3  
4 1) Digital contracts to define and enforce particular transactions and long term  
5 business relationships.  
6
- 7 2) An SDI infrastructure to suggest particular opportunities for a vendor, based on  
8 the benefits to the complete supply chain. SDI can utilize as much information as  
9 the vendor is willing to entrust in SDI, to try to identify present and predictive  
10 future benefits resulting from actions on the part of other vendors.  
11
- 12 3) A means for enabling customers and/or vendors to form coalitions, to strengthen  
13 their ability to provide pressures on other vendors.  
14
- 15 4) A mechanism with useful incentives for these customers and vendors guarantee  
16 direct benefit even if the desired action ultimately does not result in new value,  
17 e.g., an insurance scheme.  
18
- 19 5) A method to exchange terms and conditions between members in a supply chain,  
20 by being passed up the supply chain from the lowest level (consumer), and a  
21 mechanism to provide "fair compensation" for this process.  
22
- 23 6) A method to integrate loyalty benefits and long-term purchasing contracts, etc.  
24 into an electronic cash protocol, so that buyers are aware of all contractual  
25 constraints during negotiations with agents.  
26
- 27 7) Methods within SDI to advise a customer about an appropriate bidding strategy,  
28 based on estimated values of different contracts and supply chains, to help the  
29 customer to make a good decision.  
30

1 8) Methods within SDI to recommend information to reveal to a supply chain, to  
2 maximize the value of a loyalty bonus offered and ultimately received by an  
3 agent. Of course, information cannot be falsified, but some information can be  
4 withheld.

5

6 Applications of “bribing” the consumers of a vendor:

7

8 • A commercial entity is sustaining bad publicity from an article published in  
9 newspaper X. The company could offer the customer base of the publication  
10 (identified via SDI) an incentive to temporarily disrupt or boycott the publication  
11 until corrections/changes are made to the article. The company could also offer  
12 discounted products or services to the consumers, based on information (from  
13 SDI) about their consumption patterns. A more indirect discount could be  
14 provided via other vendors further up the supply chain, via multi-vendor  
15 exchange.

16

17 • Company A is a high tech start up selling software products and services to  
18 established vendor B. Vendor B identifies the substantial value proposition of the  
19 products and services provided by the start up, and decides that it should provide  
20 the services for itself “in-house”. Start up company A can go to the consumers of  
21 vendor B (information from SDI) and offer a percentage upside, e.g. percentage  
22 equity in the start up, if they will apply a threat to boycott the vendor unless the  
23 vendor agrees to do business with the startup. Company A might even identify  
24 consumers or suppliers further down the supply chain, and apply indirect pressure  
25 to B.

26

27 • Company A is developing a new commercial initiative which may be competitive  
28 to that of another smaller company B, and might harm B’s market share. Assume  
29 that consumers prefer B’s products to the future products of A, and that B can  
30 make more profits than A. Company B could convince the customers of company  
31 A of these facts, offering them a percentage of potential upside and/or discounted

or value credit towards products. The customer profiles of company A's customers also reveal the vendors from which they most frequently purchase, and company B could also offer equity or otherwise some potential upside to these vendors in exchange for offering discounts or credits to their customers (which are customer A's customers as well).

In place of monetary compensation, we might provide a value payment to an agent in the form of a long-term compensation, based on the long-term value achieved by an agent further up the supply chain.

- SDI can identify the paths in the supply chain between companies, to allow a sequential transfer of benefits down a chain. Ideally, each entity providing the benefit is doing business with another next lowest in the supply chain, so that value can be exchanged.
- SDI might "bribe" each party within a value chain to enter into a contractual relation which benefits that party (directly or indirectly).
- SDI can determine what companies along this chain would have the greatest future potential to benefit one another, i.e., where the supplier can generate compelling long term benefit to each of his subsequent buyers at each level down the chain.
- SDI can use available information regarding a vendor's internal and external activities and goals, to identify optimal multi-vendor supply chain models, in terms of maximizing the sum value to all members in the supply chain, *and* the value of each individual company.
- SDI must also consider the willingness of each vendor to establish the necessary contractual relationships for the new supply chain. Vendors might provide SDI with "relationship policies" and "contract policies", to allow this to be factored into an SDI recommendation.
- SDI models, in detail, the economic value which each scenario provides to the collective vendors and to each individual vendor. SDI can inform each individual



1 vendors of its own predicted value, while withholding information about other  
2 vendors.

- 3 • SDI can seek coalitions of vendors with similar interests to combine their  
4 bargaining power, and attempt to compel vendors to enter particular preferred  
5 value chain scenarios.

6

7 SDI is privy to most or all of the information relating to all of the companies in a supply  
8 chain, and can be trusted with the task of creating strategic plans between different  
9 companies, potentially making the supply-chain more efficient and benefiting all vendors.

10 The strategic plan might suggest:

11

12 1) New business relationships between the most synergistic entities

13

14 2) Contracts for existing commercial relationships which embody and enforce  
15 specific terms of those relationships within the framework of this optimal supply  
16 chain system.

17

18 SDI might also create a detailed system for accounting for the effect of a contract on the  
19 rest of the supply chain (individually and collectively), and predict short and long-term  
20 effects of contracts in advance. This information can be used to encourage the desired  
21 contractual activities on the part of vendors in a chain. There may also be cases in which  
22 the direct/indirect economic impact of a contract upon a vendor results from an additional  
23 secondary economic advantage to that vendor besides the trickle down economic effect;  
24 for example because of strategic commercial benefits. These factors can also be  
25 considered in an economic benefit model.

26

27 Given models to compute the comparative economic value of proposed contracts to  
28 vendors in a supply chain we can demonstrate a “trickle down” economic value chain in  
29 which the success of company A is directly related to the success of company B. This  
30 information can be presented to company A and each intervening supplier. SDI can then  
31 establish long term contracts between each entity in the chain to insure that each



1 optimal global scenario, and the optimal scenario for each individual agent. Then we  
2 could allow agents to negotiate, to find a *multi-agent compromise outcome* that is more  
3 efficient than the current solution but possibly not as good as the globally optimal new  
4 supply chain. The decision of an individual agent to enter into a proposed contract will be  
5 determined by the cumulative value, short term and/or long term, to that entity, consisting  
6 of the predicted value of taking an action (e.g. breaking a contract, initializing a new  
7 contract) both in terms of immediate bonus to a vendor, long-term individual benefit, and  
8 trickle-down benefit because of global supply chain improvements.

9  
10 Each vendor must provide SDI with as much information as possible about the particular  
11 conditions under which it will perform new contracts, such as its preference between  
12 short term and long term predicted benefits, etc. SDI looks for outcomes that are  
13 favorable to individuals and to the overall system. Value can be realized in short-term  
14 payments and long-term bonuses.

15  
16 *Additional examples include:*

- 17
- 18 • Company A may threaten to file a lawsuit against company B (which could be  
19 very damaging to company B). In a similar barter exchange, company B can  
20 **barter** its own value (in the form of upside or products/services) as offered to the  
21 customers (or vendors of the customers) of company A, targeted to only those  
22 customers that reside within the same supply chain. Thus, by limiting the value  
23 proposition to only these entities, and not to those which reside within the supply  
24 chain of a competitor, the value of the potential upside is worth more because it is  
25 partially predicated upon their own direct success. Again, it is also possible to use  
26 an indirect strategy, bribing entities or individuals of influence within a given  
27 company from which a particular action is desired.
  - 28
  - 29 • In another novel application there are certain activities among certain individuals  
30 which are considered to be mutually and collectively advantageous or  
31 disadvantageous to all members of a certain vendor community. We can have all



contract all the relevant strategic partners before a funding commitment, to provide more information and reduce risk. Other equity shares could flow from marketing agreements, e.g. a co-branded affiliation of all further products/services. The new company might also guarantee an exclusive contract with another entity. An entrepreneur might also be willing to sell a company to the more established entity, e.g., within a year or two, this can be contracted at the outset. The present framework allows value from long-term strategic alliances to be released.

Within SDI we can adapt a user's personalized portal browser to favor the vendors providing such value opportunities to the consumer. A vendor could achieve additional market share by complying with requests of another vendor. The value proposition as presented to the user can leverage the personalized browser user interface to prioritize a vendor's offers in the future.

*Example Commercial Application.*

An investor has invested in a very high risk venture relating to highly speculative stem-cell research. After years and millions of dollars, the venture finally goes out of business (or becomes acquired on the verge of bankruptcy) the investor loses all of her investment, however terms in the original investment contract (which involved at the time the approval of a variety of medical organizations most likely to benefit from the eventual medical science benefits of the research) now enable the investor to acquire in kind benefits (which may include anything from equity in these various medical organizations to goods/services provided by them). A portion of the value may even include a sizable amount of value provided by the original company and/or indirectly vis-a-vie the remaining medical entities, a substantial amount of value within the barter exchange for favors.

With time, stem cell medical applications proliferate thus, the "upside benefit" increases in proportion over time. However, if this "up-side benefit" (of the original collection of

1 contracts from the various relevant companies in the field to which the original research  
2 activities pertained) was insured through the futures market the investor can benefit from  
3 the value and growth. Investors can include researchers whose contribution of time and  
4 effort and dedicated focus on the problem had been estimated. This can allow valuable  
5 talent to receive the incentives to join the initiative from the outset in a way which would  
6 be very difficult to achieve with other standard recruiting approaches. It is important in  
7 this scenario, however not to over incentivize the managers and executives such that they  
8 may lose incentive to make the company succeed. In this scenario that the company does  
9 not go bankrupt, proper incentives could be provided to further encourage similar  
10 scientific and innovative entrepreneurial initiatives in the future, thus a portion (perhaps a  
11 fraction of the upside in these other entities could be provided even in the event of  
12 success (and irrespective of the level of risk associated with the original venture).  
13 Certainly the field of stem cell research is only one example of many speculative new  
14 technology fields with potential for enormous overall impact upon all aspects of the  
15 economy and society. Another example cited elucidates other potential features and  
16 variations of the present novel scheme. The field of na---- technology is believed by  
17 many to become perhaps the most important and pervasive technology paradigm of the  
18 twenty-first century impacting almost every industrial sector. As such, many different  
19 critical problems will need to be solved, each with potentially enormous associated  
20 economic opportunity. There is, however, considerable risk associated with the  
21 probability for industry and research institutions to realize these significant returns on  
22 investment for any given industrial sector or application domain (however, by no means  
23 is this true for the overall field of nanotech). Also breakthroughs in any one application  
24 domain will substantially advance the state of the art for potentially all other application  
25 domains. It is reasonable for SDI to create a futures market in which it predicts certain  
26 long term but very realistic goals for each application domain of nanotech. Through  
27 cooperation with industry and the investment community it may further pre-contract with  
28 various present and potentially future entities which are most likely to be direct  
29 beneficiaries of the technology which is spawned from the initial extremely high risk  
30 commercial ventures pioneering the basic technology. As in the stem cell commercial  
31 example percentages in the various longer term beneficiaries can be used as a strong



1 emerging opportunities could be disclosed as potential (suggested extension of existing  
2 departments or divisions based upon their individual areas of expertise and particular skill  
3 sets of its employees.

4  
5 The present system methodology (of business to business and business to consumer) may  
6 also be extended to include other variations including consumer to business or consumer  
7 to consumer. An example of the former could be an employee who is slated for lay-off or  
8 firing within a large organization. The indirect effect of the lay-off could result in a  
9 negative impact upon the individual's financial welfare and indirectly upon the  
10 professional contributions of that employee's children. The likely potential beneficiaries  
11 (corporations) may identify potential lost revenue streams and effectively bribe the  
12 present organization with "value" to avoid the lay off. It is also noteworthy to mention  
13 that the justification for accepting the offer for the second organization is more  
14 compelling if there exists a "trickle down" economic impact, ultimately upon the first  
15 organization (and certainly SDI, within the present framework, would factor such small  
16 economic negative factors into the optimization model for the most efficient supply chain  
17 of vendor relationships).

18  
19 Within the scope of the present example, it is possible to create a form of insurance in  
20 which SDI identifies extremely similar circumstances where similarity is measured  
21 primarily by the nature and scope of the opportunity and its associated risks as well as the  
22 entity whose critical action is necessary to bring about the desired objective. In this case,  
23 the individuals (or entities) benefiting from the desired action are able to effectively  
24 conceive to contribute a certain level of value to the collective group to effectively insure  
25 the desirable outcome on behalf of all of the members of that group, such that, if the  
26 desirable action on behalf of the other entity (in this case the large organization  
27 contemplating substantial layoffs) does not concede to the desired action for one or more  
28 of the individuals, this value contributed from the group is used to apply still greater  
29 economic pressure upon the entity to perform the desired action. If this does not facilitate  
30 the action, the value is instead applied in another (albeit less optimal) form to compensate  
31 at least a portion of negative economic impact upon that individuals/entity of course, the



1 level of risk may influence the amount of insurance each individual/entity is willing to  
2 contribute within a given group.

3  
4 1. There are numerous other extensions and potentially practical applications of the  
5 present novel methodology, particularly those applied primarily at the consumer  
6 level, and where consumers are provided with an infomediary such as SDI to  
7 represent the most individually beneficial actions within the present agent mediated  
8 contract negotiation framework. One of the significant technical advances of SDI is  
9 the ability to not only represent the particular best interests of an individual (or other  
10 entity but also provide; (1) knowledgeable experts (or "knowledge proxies") who  
11 within the information rich economic and political universe is able to off load the task  
12 of acquiring constantly updated and changing and expanding expert knowledge  
13 within each domain and sector of the economic and political systems. In this regard  
14 SDI may, as a service to consumers, represent consumers according to which overall  
15 economic scenario which best represents their socio-economic interests which may  
16 factor in both long term or short term monetary benefits and/or socio-economic  
17 benefits and/or (if they differ in any way from above) personal or political views or  
18 objectives. The use of collaborative filtering may provide a useful tool for assuring  
19 that proper recommendations for political views reflects economic political and socio-  
20 political information available about the user, such economic data of which is  
21 typically gleaned passively and may be confirmed by active confirmation of the user  
22 along with political views and positions on political issues which may prompt the  
23 user in order to collect rapid profiling information about the user is overall political  
24 objectives. Thus SDI can effectively in practical terms provide a "continual"  
25 personal advisor to guide all consumer activities (via contracts) as political decision  
26 making policies, and even these consumer actions to indirectly (via the supply chain)  
27 apply political pressure upon organizations in order to achieve certain desired  
28 economic, or even politically advantageous objectives, at an individual level.  
29 Accordingly, once they are certified on an individual level, SDI could provide  
30 synergistic benefit to the ultimate political objectives of each person individually by  
31 creating a "global political support strategy" for each group of individuals sharing the



002207-866960

feedback to politicians with sufficient notice with which to pre-formulate their political strategies at both the micro and macro levels with which to best comply with the adopted or otherwise recommended political and/or economic edicts according to SDI. Or, if it is determined that such pre-knowledge could be abused in such a way as to politically manipulate campaigns, large employers, etc. undermine the political support for these edicts, SDI may determine that it is advantageous to not disclose such information in advance to the target political groups (this, however, is a very subjective and complex analytical process). Thus SDI, representing each common group, must determine and predict which politicians are most likely to ultimately support the issues which are most important to their political objectives for which advanced disclosure of their own political support strategy is likely to advance the strategic knowledge of that politician while assuring that none of their positions on none of the political issues are likely to be ultimately abused by that politician by him/her ultimately changing his/her position on that issue in order to improve his/her chances of election. For this reason, it is also of interest to each group to attempt to predict the ultimate position that each other group is likely to assure for each political issue in order to determine whether or not pooling their collective influence on that particular issue (including free disclosure to those "trusted" candidates) would be advantageous or information which would be potentially harmful to pre-release to a potential adversary on that issue (which it may ultimately disclose to a candidate opposing that issue and/or modify its own political strategy in a manner which is ultimately antagonistic to the interests of the welfare of that issue or theoretically even issues which that group may wish to disclose in the interest of providing further support for the global political strategy through this form of coalition creation). It is worth noting that within the scope of the present framework it is likely that the availability of such detailed information, both predictively and at an individual level, regarding the ultimate political and economic effects which certain voting in other political support actions via SDI are likely to imply, it is extremely likely that members of each group may be reluctant to provide information, regarding the groups political strategy to public polls.

Thus it is a primary objective within the scope of the present system's formation of common political support strategy groups to ascertain each individual's user's ultimate political objectives based upon the analysis which provides analysis and prediction of the present and likely future effects, both politically and economically from the standpoint of that individual (which may include social, professional and predicted individual investor oriented direct and indirect consequences at a detailed level as it relates to that individual. And secondly, to formulate a global political strategy based upon these objectives which are represented by an SDI agent and which is able to recommend actions and, as such, act as a coalition on behalf of the constituent individuals, and including negotiations for further coalition formation on various portions this global strategy in order to further enhances the collective power wherever common interests can be safely identified and shared to both groups' mutual best interest.

The effects on commercial industry, even specific businesses in which certain actions are likely to result, may be modeled and presented to the user as well as the effects upon the user in light of the investment portfolio of the user. If a particular political position were to be taken, the system could even recommend-----  
-----which is most compatible with a particular political position.

## Continuation to Supply Chain Section

It is also possible to utilize the extensive information in SDI to instruct each political support group's most advantageous supply chain structure in order to be able to recommend certain strategically advantageous cases of tax dollars which facilitate the construction of such supply chains in accordance with the above methods above described. It should be noted that a "political group" may be further subdivided into groups based upon purchase affinities (which themselves suggest different group divisions) which in turn may bribe or boycott commercial entities as needed to achieve their most advantageous strategy as consumers or simply individual users (as economic entities). Each subgroup (via SDI) may further augment their economic leverage with

1 their political leverage (over tax allocation decisions). E.g., if candidate X supports this  
2 particular supply chain strategy and if he is elected then it is most optimally advantageous  
3 to also utilize consumer leverage in supporting transactions with certain vendors and  
4 boycotting others in order to support or complement the politically (tax) funded strategy.  
5

6 As suggested, SDI may be used to facilitate the matching and strategic formation of  
7 coalitions of groups which are most likely to benefit one another on particular political  
8 issues or formation of strategic contractual relationships between commercial entities  
9 (SDI is able to determine if/when certain information is helpful or harmful to disclose to  
10 another particular entity). Each group may exert particular leverage over politicians. I.e.,  
11 not only could the politicians receive feedback as to the political issues and economic  
12 (commercial) alliances which are most advantageous but also the groups could upon the  
13 advice and direction of SDI make their votes conditional upon the politician supporting  
14 certain issues or economic actions. SDI could provide real-time feedback right up until  
15 the election as to the numbers and percentages of users who had supported the candidate  
16 and the candidates response to them. There are other instances in which the leverage  
17 could be used to incentivize commercial entities to support certain political concerns e.g.,  
18 if company X donates percentage of its revenues to the United Negro College Fund group  
19 Z will support the facilitation or an alliance between company and a major manufacturing  
20 firm (e.g., by putting consumer pressure upon those retailers which are the  
21 manufacturers' customers to in turn threaten temporary boycott of the firm unless the  
22 stipulations are met or if they are not consumers to instead offer consumer loyalty (e.g.  
23 through dedicated currency).  
24  
25

26 A variety of applications of the present scheme are conceivable such as local politics such  
27 as commissioners office and board members delegated for purposes of masterplaning and  
28 use applications and allowing or disallowing certain real-estate and highway development  
29 projects. Certain commercial projects, for example, could be of substantial benefit to  
30 certain commercial entities which could provide considerable direct long term economic  
31 benefit to citizens in the local community even directly to the individuals in the

1 communities (who elect those officials) as well as contribute indirectly via the available  
2 tax base. Or certain commercial entities could, for example, directly or indirectly provide  
3 economic benefit to certain individuals, e.g., if they are employees, investors or  
4 commercial entity which is in some way commercially dependent upon that company,  
5 e.g., for commercial business if it is a customer or alternatively a supplier and the impact  
6 of its resulting success in general (or even local presence) would again provide trickle-  
7 down economic benefits to that entity. In accordance with the present application it may  
8 even be possible to estimate values for such factors as, for example, the economic impact  
9 20 years in the future upon commercial business and industries in the local region (as  
10 well as nationally) if local schools are provided with a 10% vs. a 7 ½% share of the local  
11 tax base. With these predictive models, the associated prospectively affected commercial  
12 entities could, effectively encourage SDI acting on behalf of the local citizens to form a  
13 coalition which makes their voting of the relevant elected official(s) contingent upon the  
14 increased tax allocation for local schools (the specific preferred users could be predefined  
15 as well). SDI acting on behalf of the interests of the prospectively affected commercial  
16 entities could bribe the citizens with either (preferably) such benefits as in kind,  
17 goods/services which could in turn be leveraged through the supply chain or through a  
18 barter exchange system or employment benefits or wage increases (which again could be  
19 effectively treated as a barterable commodity subject to the above conditions of avoiding  
20 interactions with competitive entities. Such modelling could be extended to many other  
21 situations, e.g., allocation of funds to police, fire and EMS services, approving and  
22 providing highway improvement again may provide the above benefits etc. providing  
23 safety and preserving the health, welfare and life (in addition to educational quality) of a  
24 certain number of individuals (each with a certain predictively modeled) local economic  
25 impact. The cost versus impact affecting each citizen could be presented to them by SDI  
26 accordingly in order to elucidate their decision making processes.

27  
28 There may be certain instances in which users at an individual level may be able to provide  
29 predicted -----this benefit can be quantified towards commercial entities or  
30 even individual users if particular events can be achieved (or particular events prevented).  
31 For example, consider a student who is determined to be capable towards achieving



Information which is enclosed may, for example, contain in addition to the identity of the prospective investors, the complete discussion of the economic benefits which each individual investor could potentially sustain, however, this information is disclosed only to the SDI representing the collective interests of the prospective investor coalition (a primary directive of which is to never compromise the data disclosure policies of any individual investment entity, even if such disclosure to one or more of members of the group may ostensibly benefit the collective group). It is thus the objective of this collective SDI agent to make individual data disclosure recommendations for each prospective investor that will optimize the probability that the basic framework of the collective cooperative group is successfully created while enabling the individual entities to withheld sufficient strategic information from the group in order to optimize its own negotiating position with the other entities. For example, it would be particularly advantageous to disclose credentialed information regarding the benefits which the investor could provide to the collective group of investors (based upon the existing investor information disclosed by the associated companies) and information regarding the existing relationships which that investor has with existing companies which could, in turn, provide positive benefits to the companies constituting the global investment strategy of the group and (in the case of negotiating with each given individual investor prospect) the benefit to companies which that investor has a relationship with and made available to the present investor prospect. These companies may also include those which SDI has recommended to that investor as part of the global investment strategy for that group.

user's reactions to various events and stimulate the video programming can be captured, aggregated and accessed by present viewers such video in real-time or asynchronously for future viewers of non-live content by user profile (or content profile) or by similar users to that of the profile of the user. Observing reactions/responses to political or ethnic jokes by that group which is targeted or other groups may be interesting to users to observe as a simple example. Reactions/responses to news or political events by different groups or those most effected thereby may be of interest to users.

## 1.7 Off-line Variations





1 If a smart card is used this user profile data may not have to be remotely retrieved but may  
2 be stored on local memory on the card itself along with the client-based pseudonym proxy  
3 server. In one novel variation, a card is done away with completely by virtue of  
4 revolutionary technological breakthroughs in being able to instantly and positively identify  
5 buyers biometrically using iris scanning techniques (which may in a variation be further  
6 combined with facial recognition techniques). Many vendors will wish to utilize user profile  
7 data in order to deliver promotions targeted discounts and promotions (see pending patent  
8 “System for Customized Prices and Promotions”)

### 9 **1.7.1 Location Enhanced SDI System**

10  
11 The co-pending application entitled “Location Enhanced Information Architecture” (LEIA)  
12 describes an integrated advertising delivery platform which selectively targets buyer  
13 personalized advertising based upon both the buyer’s personal profile and the present  
14 location of the buyer which may suggest appropriate ads from vendors which are local to the  
15 buyer, wherein user identifiers (UID’s) which could include any of the above identification  
16 media provide the essential elements for this buyer targeting platform. With LEIA  
17 information providers can utilize location information, in addition to static and dynamic  
18 profiling information. The method customizes the information that is displayed on a private  
19 or public information device to the real audience in the vicinity of the device, instead of a  
20 predicted audience. LEIA collects an extremely detailed and comprehensive information set  
21 about the daily activities of a user, enabling enhancement of the user profile with location  
22 information and temporal activity patterns. The co-pending LEIA patent suggests  
23 appropriate application environments, for example in a smart home, an office, on a mobile  
24 shopping device, and in an automobile. A LEIA-based system stores personal information  
25 on users.

26 We can extend LEIA by incorporation with the Secure Data Interchange system that we  
27 teach in this patent. SDI enables the user to receive the benefits of powerful and well-  
28 directed information, but within a system that respects his/her privacy requirements. The  
29 interchange acts as a secure data warehouse for users and information providers, enabling  
30 information providers to target users without revealing private information to the providers  
31 directly. As described in LEIA, the privacy architectures provided for in LEIA and SDI are





Other extensions of this scheme are also considered e.g. within the context of the user's office, or automobile and pedestrian activities. This application may thus extended the usefulness of the iamworthit model to advertisers in being able to target users through the presently anticipated on-line media as well as networked appliances and in either case, based upon the relevant context of users' present activities and behavior (and from this potentially their inferred moods or mental states) within their homes and elsewhere. Additionally (perhaps most importantly), LEIA provides a highly beneficial value to vendors whose customers purchase primarily from the vendor's bricks and mortar store presence. SDI can identify multiple vendors which share a common physical location. Additionally, LEIA can also identify the physical location of buyers on mobile and terrestrial networks and suggest the sharing of off-line prospects for these off-line retailers on the basis of not only interest but physical location.

### **1.7.2 Digital Set Top Box Methods**

Similarly, we can deliver targeted advertising and other information through cable TV systems, as described in the issued parent patent application entitled "System and Method for Scheduling Broadcast of and Access to Video Programs and Other Data Using Customer Profiles" US Patent # 5,758,257, and co-pending application entitled "Broadcast & System for reduced memory terminals broadly address the use of cable systems as an interactive medium (in a bi-directional network architecture) for purposes of delivering targeted advertising targeted advertising and other information to the consumer based on user profiles".

In this system customer behavioral data is collected at the digital set top and the upstream channel enables these profiles to be processed at the lead end server. These detailed profiles may then be subsequently transmitted down and stored at the level of the individual set top. The cable environment is a two way interactive medium. The bandwidth allocation is inherently asymmetric. Separate channels can push parallel adverts, which are selected at the set-top-box according to a buyer's profile. Each channel can have associated meta-data to allow matching at the set-top-box. As an alternative variation, full motion advertisements may be down loaded in the form of applets to the digital set top box and displayed to the

1 buyer in similar fashion as described above. Or full-size commercials or infomercials could  
2 even be downloaded and inserted in place of existing commercials on the video stream.  
3 Digital tags used for queuing for ad insertion technology enable this capability. The  
4 preferred commercial implementation of the system is used within the context iamworthit  
5 (the buyer infomediary service in which the viewer receives value for advertising received).  
6 This method is preferred (in as much as in the following alternative, the cable operator does  
7 not need to be in the loop). With that said, in an alternative embodiment, however, requiring  
8 the cooperation of the cable operator various commercials could be broadcast in parallel  
9 during commercial breaks. The user profile at the set-top is deployed to select the most  
10 appropriate commercial accordingly.

11 User profiles regarding viewing behavior can be collected with relatively minor up-stream  
12 transmission, e.g., to infer whether a buyer is viewing a program the set-top may transmit  
13 the viewers current viewing selection, e.g., two minutes after the beginning of each half hour  
14 and two minutes before the end of the half hour (and possibly at additional intervals during  
15 the viewing segment as well). In the point-to-point access and delivery of personal  
16 information as well as the passive monitoring of viewing behavior (for profile generation at  
17 the head-end and subsequent user profile delivery for the set-top) because of strict buyer  
18 privacy regulations in the cable/satellite industries the use of the pseudonym proxy server  
19 architecture (as described in the parent case and integrated into the above referenced case) is  
20 extremely important.

21 Once interactive television is enabled by access to very large amounts of bandwidth, it will  
22 be very feasible from a bandwidth utilization standpoint to transmit more interactive  
23 content. Predictive caching is still extremely advantageous in as much as advertising (unlike  
24 interactive video real-time on demand video content) tends to be controlled by targeting  
25 rules of the advertiser rather than the user who receives value in exchange. Of course, radio  
26 or music content may also be delivered in conjunction with preloaded audio commercials.  
27 And any of the above content is deliverable over the same networks to other appliances such  
28 as the P.C.

### 29 **1.7.3 Optical-Based System Extension**

30 Iamworthit in a future version of the system could incorporate a novel optically based  
31 medium for delivery of the advertising. The protocol would enable the delivery of user

targeted advertising in a non-electronic environment directly to the user from a static source in visible contact of the user. The system incorporates the following aspects: The user wears glasses or contact lenses which contain a Polaroid film in which the optical medium is polarized for only certain very narrow and specific wavelengths. Advertisements (or other messages) may be presented from signboards which display from the same surface ads appropriate for every different user profile segment for which a unique message is appropriate. The Polaroid film for each user is unique to the user segment to which they belong. Each user segment's optical film filters out all of the particular wavelengths of all other messages except that which is appropriate to the user segment to which that user belongs. Alternatively, the current application may also be relevant to virtual tags as users may provide instead optical messages specifically targeted by a user to other users characterized by their profile features which establishes the entire criteria for the user segment they fall into for which an appropriate message may be targeted.

## 2. Interactive Data Mining Applications

In this section we describe interactive data mining applications, which involve dynamic two- or multi-way communication between agents. For example, within the context of a vendor-consumer interaction, the key difference is that information flow is not one-directional. An important example is time-of-purchase, in which the consumer's agent initiates a request for a response (with a counter bid) from a number of vendors. The vendors can access profile information about that consumer, and then make personalized responses. In general, this bid-response processes can iterate, with vendors/consumers exchanging information over multiple rounds. The exchange of information may occur within SDI in initial stages, with the system acting as a trusted intermediary until the agents are ready for an introduction.

We describe time-of-purchase competition, and then move on to general match-making applications, which allow interested parties to exchange information by mutual consent (as indicated in SDI rules), with initial information exchange autonomous. There are a number of interesting match-making applications, for example within a system for sensitive negotiation, introductions. The system of SDI can play the role of a trusted intermediary, so that only suitable introductions are

made. A vendor can place a request for a certain type of meeting without its competitors knowing that it has made such a request. Finally, we describe a “real-time experts market”, in which experts can respond for payment to questions placed by other agents. The system of SDI allows a useful matching of agents and experts, and also allows data mining to check on the feedback a particular expert has received from previous clients. We also suggest an application of interactive SDI-based data-mining to resale markets, and to a decentralized transportation management system.

## 2.1 Time of Purchase Competition

One application of the system of Secure Data Interchange is in “time of purchase” competition, where by a networked system of vendors can subscribe and receive the opportunity to place counteroffers to users that are about to make a purchase. This application is enabled within SDI because the user agent can remain anonymous while receiving counteroffers, but still use the system of SDI to provide profile information that can allow vendors to make attractive personalized offers. The client-side SDI proxy notifies a central SDI “time of purchase” server, which can:

- Determine appropriate vendors to prompt with an opportunity to make a counteroffer.
- Create a one-time anonymous identifier for the user, linked to an appropriate set of profile information, and allow vendors to execute queries under that profile.
- Collect counteroffers, and pass them to the user’s client machine.

We also describe how to generate coalitions dynamically, based on the privileged position of the time-of-purchase server in the marketplace. The time-of-purchase server can transparently bundle purchases from disparate end-users, and negotiate quantity discounts with vendors. Buyer Coalition formation can significantly improve negotiating leverage as the collective buying power of numerous individuals using the time-of-purchase competition feature, and be made completely transparent to users.

Time-of-purchase allows a user to elicit dynamic market competition between vendors. For example, when a user browses a book at Amazon.com if she/he chooses to activate the time-of-purchase competition feature from the browser, a notification will be sent-out in real-time to all competing book stores which also offer that product section. The notification contains





1       b) The vendor is un-cooperative, and the client-side SDI proxy needs to use other  
2       techniques to classify the type of product or service that a user is requesting from  
3       a vendor.  
4

5       The central SDI data warehouse maintains a searchable index of vendors that provide  
6       certain products and services, and sending buyer-profile and purchase information to  
7       appropriate vendors just before a purchase is made. Vendors can register with SDI to receive  
8       time-of-purchase competition opportunities, and can also register to allow time-of-purchase  
9       competition with purchases on their web pages (case a, above). Of course, it is unlikely that  
10      a vendor would do this without incentive, perhaps one incentive could be a discounted price  
11      for registering to receive opportunities for time-of-purchase competition.

12     In fact, this problem of client-side identification of transactions is a key problem that must  
13     be addressed in a description of client-side user profiling and the submission of information  
14     to SDI. The system of SDI requires a method to know what the user is doing. Click streams,  
15     typed query information, provided profile information etc.

16     A vendor that registers to receive time-of-purchase competition opportunities agrees on a  
17     standard classification system for products and services, and to state what types of products  
18     and services it wishes to enable for time-of-purchase competition. The classification system  
19     encodes the product or service that a user requests, at an appropriate level of detail to allow  
20     other vendors to make reasonable counteroffers.

21     The central server maintains a database of vendor interest sets, so that it can  
22     determine from a classification code which vendors will be interested to provide  
23     counteroffers.

24     For example, if the product is "flights to the UK" then the code might encode the dates that  
25     the user wants to travel, and his/her preferred departure and arrival airports.

26     Take case (a) , where we assume that a vendor is cooperative and provides the SDI proxy  
27     agent that implements time-of-purchase for the user with a final product/service code and  
28     price. The system works as follows:

29       1. User enters into a dialogue with vendor.

30       2. Vendor makes "final offer" to a user, and the user is about to make a purchase.

31       The vendor has also informed the user of the product/service classification code





4). B2B Commerce – Numerous commercial applications including standard Web-based environments for B2B e-commerce including buyer-side and seller-side e-commerce applications, vertical B2B portals and trading floors.

5). Job Positions – Employers posting job positions may be notified by the service of competing job offers which employment candidates are browsing ( and ultimately accept if this information is indicated by the candidate). Resumes and, if desired, other profile information is provided to all competitors. Depending upon the profile of the candidate and/or competitive employment opportunities which a particular candidate is observing (or has recently observed), employers may wish to customize offers dynamically on an individual basis.

#### 2.1.4 Integration into a Shopbot Interface

We also suggest the integration of time-of-purchase competition into a standard shop-bot interface e.g. with comparative features of the products associated with the various offers across a variety of product criteria, or at the least price can be used to provide the user within simple recommendations. With this, even without a time-of-purchase offer the user can at least compare its offer with fixed-price deals in the wider market place. We might also provide a user with historical information about previous (low) price offers, within a Yahoo-like portal for e-commerce; to provide a user within information about a reasonable price for each individual purchasable; i.e. the very lowest price that each given item had been offered to a customer previously using time of purchase competition. This site could also act as a stand alone proxy server (like anonymizer) which inserts this information by overlaying it on Web pages as the user browses the web.

The time of purchase competition feature is designed to be accessible to a user within the context of or most any information access mode relating to browsing on the Internet e.g. Web or portal browsing, receiving email or “push” content, submitting queries (e.g. specific names, categories and/or desired features (or combinations of the above) of products which the user identified as being of present interest.

1 During any of the above information access modes when a buyer specifies *time-of-purchase*  
2 *competition* at her SDI client proxy, SDI can automatically provide competitors with  
3 information about a buyer's product or service requirements, and a buyer's profile, as well  
4 as on particular vendor specific offers the buyer is observing or contemplating before a  
5 buyer makes a purchase. This can facilitate competition between vendors, and can lead to  
6 better prices and offers for buyers. The system of time-of-purchase competition can also  
7 help to reduce the costs of entry into a market, because name-recognition becomes less  
8 important. New vendors can simply register with *time-of-purchase referral system*, and  
9 cherry pick a small set of purchases.

10 This is a next-generation e-commerce service. Current shop-bots, for example "Junglee" at  
11 Amazon.com and [www.shoptheweb.amazon.com](http://www.shoptheweb.amazon.com) provide a *static* comparison shopping  
12 service. A buyer can specify a product, and receive price information about the product from  
13 different suppliers. There is no dynamic competition between vendors on price or features.  
14 The buyer driven service for flights offered by [www.priceline.com](http://www.priceline.com) is more dynamic, in that  
15 a seller is found to match the price that a buyer bids, but does not necessarily promote  
16 competition between sellers for a user's purchase. In fact the sellers can make excess profits  
17 from the pricing errors made by buyers, because once a buyer states his/her reservation  
18 price, that is the price he/she must pay. Surplus goes to the market operator, e.g. priceline,  
19 but not to the user. The system is also different to our proposed system in that:

- 20 a) Time-of-purchase competition is disintermediated and decentralized,  
21 implemented dynamically over a virtual network of competing vendors.
- 22 b) Time-of-purchase also allows a vendor to access anonymous profile information  
23 about a buyer, to enable a careful tailoring of an offer to the preferences of a user.

24  
25 We can use profile information, and historical transaction information for similar  
26 transactions, together with the customer price/promotion algorithm disclosed in co-pending  
27 patent "System for the Automatic Determination of Customized Prices and Promotions" to  
28 negotiate on a deal with a vendor that will optimize the value to the buyer. Profiling of  
29 vendors, and buyer transactions, can allow buyers to avoid making bids that are too high and  
30 losing value (airlines in priceline.com can profit from inaccurate buyer bids).

1 Time-of-purchase enables vendors with competing products or services to receive automatic  
2 notification when a buyer is about to purchase a relevant product or service. A vendor can  
3 also receive information on the profile of a buyer, and the offers made by other vendors; and  
4 submit counter-offers to a buyer via the buyer's SDI-enabled client. The buyer can then be  
5 presented with a final set of offers, before making a purchase decision.

6 A buyer can also configure its profile management rules within time-of-purchase to provide  
7 profile information relating to the buyer's sensitivity to discount offers, customer loyalty  
8 with other vendors, value responsiveness (bargain driven), responsiveness to high quantity  
9 discounts (for only those categories which the buyer makes frequent or large purchases), etc.  
10 This information can be *certified* by the buyer's client-side SDI proxy, as a fair  
11 representation of the purchasing habits of a buyer, and can be aggregated across product  
12 types to protect a buyer's privacy.

13 The SDI time-of-purchase server can identify vendors with similar products or services,  
14 either using a static index which maintains vendors in particular product domains, or  
15 through dynamic profile matches between the target object profile of the web site that the  
16 buyer is currently browsing and target object profiles of the web sites of other SDI-enabled  
17 vendors. Alternatively, perhaps classification and clustering techniques can be deployed to  
18 identify similarity between vendors at the level of target objects, i.e. the products that a  
19 vendor offers; as indicated by virtual tags provided by a vendor that profile its products and  
20 services.

21 The time-of-purchase competition system may also draw in static "listed" offers on a  
22 vendor's web site, to prevent a vendor over-charging a user. In addition, SDI might track  
23 purchases and via clustering techniques request prices on similar or competing products to  
24 the product that a user is about to purchase. These products and/or services may also be  
25 released to a user. When presenting a choice of products and/or purchases to a user we can  
26 allow many features to be summarized, for example features which relate to the location of a  
27 vendor, terms of shipment and insurance, and user ratings.

28 Vendors are notified, and provided with the ability to access the profile of the buyer (as the  
29 buyer deems appropriate), either with client-level processing or through the release of an  
30 anonymous profile to the vendor. Vendors typically will wish to construct offers through a  
31 rule-based engine, data-mining techniques, or automatic collaborative filtering techniques,

as disclosed in co-pending patent application "System for Automatic Determination of Customized Prices and Promotions" and U.S. Patent #5,754,939, "System for Generation of User profiles for a System for Customized Electronic Identification of Desirable Objects" as such techniques may be deployed by the vendor directly or via the Secure Data Interchange representing the interests of the vendors.

*Choosing an Offer to Make*

User profile information may include a temporal profile of the buyer's present activities, including search terms, recent page navigations, what pages is the buyer observing presently (and the profile of this page) or even his/her present physical location as well as the general user profile. Any portion of the above particularly the latter two may of course be withheld from the vendor). Or the buyer may be more generous in providing this information about the various static or dynamic features.

Vendors can target buyers on the basis of their preferences and interests, and also within the temporal context of when they are most likely to be receptive to offers (thus providing also a benefit/service to the buyer as well as enabling the competing vendors to increase the price point at which their competition begins, knowing that the buyer is in an optimally receptive condition to accept offers for that item). For example, a user might release to a vendor the current context of a user's purchase, such as the click stream over the past 5 minutes.

In the preferred implementation, vendors are also provided with a (client or web-based) rules interface which enables the vendors to input pre-stated rules with which the system may solicit and respond to competitive offers automatically. In this way data mining may be performed in order for the vendor to determine what a user (or all users) by attribute, tend to best respond to by product, product feature, features or services of the vendor, price, etc. If pre-stated rules are used to automatically respond to a notification with a competitive offer, the nature and degree of discount is typically determined in accordance with the nature and degree of the original or previous offer and/or the user profile as disclosed by the client-level proxy/server to that vendor. In lieu of manually entered rules, co-pending patent application entitled "System for the Automatic Determination of Customized Prices and Promotions" another similar algorithmic methodology may be used as an aid by the vendor in order to automatically determine a competitive offer (or subsequent responses thereto) as well as an aid to the vendor in selecting optimal rules.











buyers in the past, and form a model of vendors, to determine whether or not a buyer has received good offers (i.e. we can exchange information within the system of Secure Data Interchange, and making more information available increases the efficiency of the market). Offers can be filtered and presented to a buyer in rank order.

### **2.1.5 Time-of-Purchase with Non-cooperative vendors**

We suggest a simple technique to allow a client-side SDI proxy to automatically detect that a purchase is about to be authorized, and the details of the purchase. The model is that SDI, or a third-party SDI client, provides a banking service for a user. The advantage of such a service is that in making a purchase, we can require that the vendor provides information to the bank about the purchase that is about to be made, as a condition for validating payment.

The system works as follows:

1. User interacts with vendor, and decides to make a purchase.
2. User provides payment method, for example an SDI bank account number.
3. The SDI bank account server demands information about the product which is about to be purchased before authorizing payment for the service.
4. The vendor provides the SDI bank with information about the purchase.
5. The SDI time-of-purchase system can now step into the transaction, and solicit competitive offers from other vendors.
6. The user is presented with the opportunity to select an alternative offer.
7. The SDI proxy-agent reports the user's final purchase decision to the appropriate vendor, and denies purchase to the initial vendor if the user accepts an alternative offer.
8. The final transaction is executed.

### **2.1.6 Automatic Buyer Coalition Generation**

The time of purchase competition system can also be used to allow the automatic formation of user coalitions. Coalitions can be generated dynamically based on the privileged position of the time-of-purchase server in the marketplace. The time-of-purchase server can transparently bundle purchases from disparate end-users, and negotiate quantity discounts





1 utilize the above modeling techniques in order to establish pricing models based on the  
2 value of these different items in light of the total monetary value which the coalition  
3 represents to the vendor collectively. In addition to inferring buyer interest entirely  
4 passively, it is possible for buyers to manually edit their profile, e.g. by modifying or  
5 actively selecting categories of purchasables of interest which she/he would like to be  
6 notified of (e.g. via pager) if/when a buyer coalition for a purchasable of that type is being  
7 established. It is further possible to increase the size of buyer coalitions by allowing a  
8 certain amount of time for buyer's to respond to offers to join that particular coalition (e.g.  
9 24 hours). Typically a couple of days is all the time that is required to achieve a  
10 substantially maximum buyer response. Also, because some vendors may wish to be  
11 (automatically) informed if/when substantially large coalitions approach the vendor, there  
12 may be an additional marginal price advantage by providing an additional period for  
13 vendors to reply.

14 The buyer may wish to allow the iamworthit time of purchase vendor competition system to  
15 operate in fully automatic mode whereby certain "acceptance parameters" are provided to  
16 enable the buyer agent to act fully or autonomously on behalf of the buyer's stated (or  
17 approved) parameters. This technique is likely to provide a central function particularly in  
18 most business-to-business iamworthit commercial applications. The business-to-business  
19 commercial domain may exemplify a couple of other novel features (which are not  
20 exclusively limited to business-to-business). For example, because replenishment of  
21 products or supplies suggest much greater repetition and thus predictability, it is sometimes  
22 useful to buyers to also leverage as part of the negotiating process an advanced commitment  
23 of buyer loyalty over an extended period of time. In the unlikely event of the buyer  
24 changing his/her vendor loyalty after considerable value has been provided to the buyer, the  
25 vendor may control the right to acquire certain assets of the buyer (e.g. corporate stock) as  
26 pre-agreed collateral for the vendor. Also, in the business-to-business domain iamworthit  
27 acts essentially as (or in conjunction with) a buyer-side portal to a web-wide vendor  
28 resource, though it could also be integrated as part of a seller side-portal interface tool as  
29 well.

30 As an additional service to vendors SDI can provide enhanced profile information,  
31 aggregated from other vendors, to enable vendors to provide better focused offers than can



1 be provided on the basis of the profile information directly associated with the pseudonym  
2 of a buyer. Certain portions of the user profile data that is unavailable for direct collection  
3 by the vendor (such as information that is collected on other sites including, in particular,  
4 competitive vendor sites) may reveal important information which enables the vendor to  
5 better target that buyer. As such the secure data interchange representing the collective  
6 buyers may aggregate, analyze and sell this data to the vendor so long as the release of such  
7 information does not negatively affect the predicted pricing or value levels for that buyer as  
8 performed by the above type of pricing algorithm.

9 An interesting variation of this example involves the situation in which multiple entities  
10 represented by SDI users may be negotiating with the same other entity (or individual). It  
11 may be for the same purchasable. Alternatively, it may involve different purchasables or the  
12 represented group may even be a combination of users and sellers who happen to be  
13 presently dealing with the same entity. Accordingly it is important for SDI to achieve an  
14 understanding as to the individual negotiating parameters as well as an assessment of a  
15 market demand model which characterizes the needs and objectives of that entity with  
16 regards to the particular prospective transactions being negotiated which SDI is mediating  
17 on behalf of each party which is captured through the main SDI server. As such it is often  
18 possible to thus attempt to predict the minimal acceptable terms of an offer which would  
19 meet those conditions. The negotiating leverage from all parties is based upon terms which  
20 affect multiple (or perhaps all) of the deals collectively. SDI thus has the power to  
21 significantly strengthen the negotiating leverage based upon a market demand model which  
22 characterizes that entity thus benefiting the entrusted parties on the whole. The statistical  
23 data used for market demand models from the user-centric or vendor-centric SDI is  
24 collected about the other parties which the parties are negotiating with. One useful source  
25 for estimating this predicted market demand curve which they represent. For example it is  
26 possible for end-users as a coalition of users to be automatically formed for this purpose.

27 A user-centric SDI is then assigned to the coalition in order to (a) ascertain from user profile  
28 information (in addition to subsequent active query or recommended offer feed  
29 back/approval information about terms which the coalition would likely wish to receive  
30 from the vendor (subject to negotiation) (b) ascertain from vendor behavior relating to past

users and coalitions terms which would be reasonable to expect from the vendor in order to anticipate reasonable offers/counter offers to present to the vendor.

SDI's predictive tool suite can be used to predict the most likely prospects for a given coalition, i.e., matching users who have explicitly indicated an interest or have engaged in negotiations or past dealings with a given vendor. As indicated the recommended users may be targeted either with the same items or completely different items, the latter case in which SDI will try to negotiate a lower price based upon the aggregate sales from the coalition, not the volume of sales for one particular item per se, as in the first instance. In the present application, however, because negotiating pressure is being applied by a collection of entities upon (typically) one entity (or possibly alliance of entities), an SDI (which may typically be spawned from the main SDI service) representing the coalition is created in ad-hoc fashion. The other entity may accordingly wish to acquire its own SDI to also represent its commercial interests. Each SDI may represent multiple parties. No SDI must of course ever represent an entity which is negotiating with another entity which it represents.





1 these complex vendor synergies exist even indirectly via one (or multiple) removed vendor-  
2 vendor and user-vendor relationships.

3 *2.1.8.2 Coalitions Created For Purposes Of Eliciting Change On The Part Of Vendors*

4  
5 In addition to applying market pressures, one other example in which entities may be  
6 automatically introduced into coalitions using SDI may involve SDI automatically  
7 identifying certain situations or complaints from users or entities in which there may be  
8 common ground for these parties to collectively apply political or legal pressure upon  
9 another entity from which the problem arises, or in one example to apply legal tactics for  
10 purposes of reclaiming damages from the infringing party. For example, members of SDI  
11 could be asked to disclose particular instances of these types of legal infractions even if they  
12 are relatively minor (e.g. simply a short fall in acceptable quality of misrepresentations  
13 through misleading advertising etc. may be relatively minor inconveniences and/or a site  
14 may be provided as well for non-SDI members (e.g. [www.class.action.com](http://www.class.action.com)).

15 A similar application tailored specifically to the Federal Government could also be provided  
16 (e.g., [www.governmentindustries.com](http://www.governmentindustries.com). In one example, SDI's ability to provide annotations  
17 excluding complaints and ratings pertaining to particular organizations and individuals is  
18 complemented by the current scheme by introducing the ability of users to apply pressure to  
19 the organization or individual as a coalition in order to more effectively motivate the desired  
20 change accordingly. To this end other tactics may be used, for example, a threat to  
21 collectively sponsor negative advertising or editorials. In fact, the above technique for  
22 presenting a believable threat could be utilized in this case and backed up by an insurance  
23 policy which provides in addition to the guaranties that the threatened action will be carried  
24 out, for example, by X date unless a particular condition is met, also provides sufficient  
25 funds to follow-through accordingly and in one variation further instill the motivation on  
26 behalf of the parties by actually insuring the outcome of the litigation (where the threatened  
27 (and guaranteed) consequence is sufficiently great that the entity is much rather proved to  
28 settle thus reducing exposure significantly).

29 Within the present system SDI provides the framework by which appropriate users and  
30 sellers may be matched together. It also enables a methodology by which the user interests

are protected through the use of matching of sellers offers to competitive vendors (using  
iamworthit).

### **2.1.7 Single-user Time-aggregated Purchases**

Another extension is to propose a special interest-bearing bank account which we set-up for  
the user which is tied into a wallet or debit/credit card for off-line transactions. We can  
allow a user agent to bundle its own purchases over time, and use the system of SDI to  
guarantee future payment to on-line vendors in return to a good current offer. This is novel,  
because it is like participating in a “discount scheme” without purchasing the right to future  
discounts up front. A user can still invest money that is earmarked for future purchases with  
a vendor.

We describe a wallet application for this purpose, that allows a user of SDI to commit  
currency to vendors for future purchases. The time-of-purchase competition system can  
allow vendors to see the pseudonymous purchasing profiles of users (e.g., over the past 12  
or 18 months) which are relevant to their respective commercial venues and competitively  
bid for the opportunity to acquire from that user a “pre-commitment” for similar spending  
amounts within the same time period going forward (provided that that vendor offers the  
same/similar categories from which the user had purchased that amount in the past from  
another vendor(s).

The “committed currency” in the user’s account is not actually spent and thus remains  
interest-bearing for the user, and because the dollars are exclusively and irreversibly  
earmarked, the actual debit could even occur substantially after the actual purchases  
ultimately occurs with that vendor. Vendors marketing the service may be able to exclude  
competitive offers (as is the case with the standard time-of-purchase competition  
application).

### **2.1.8 Special Vendor Treatment in a Community Dollar Scheme**

1 There is a variation of community dollars in which some of the competitive features of  
2 the time of purchase competition system can be utilized by vendors to provide a  
3 competitive advantage over other vendors.

4 Each vendor is given exclusive rights to information about offers which users who  
5 subscribed through their site or portal are currently considering. All other competitive  
6 vendors are denied access to time-of-purchase competition for that user. The vendor may  
7 request profile data from the user which s/he may wish to optimize the selection/features  
8 of the particular product offering as well as the value of the offer, and use SDI or its own  
9 price engine as before.

10 If the user has also subscribed to full time-of-purchase competition, he/she may be  
11 willing to give the vendor the right to provide the first counter offer. If the user does not  
12 accept the counter offer, the system can default to the standard multi-vendor competitive  
13 scheme.

14 Vendors signed up with the community dollars scheme may also be provided with special  
15 privileges for advertising to a user as the user browses the Internet. This can be used to  
16 encourage more vendors to sign into the SDI/community-dollars scheme, to be exempted  
17 from ad replacement technology and receive other preferential treatment.

18 The present time-of-purchase competition scheme can also be adapted for the application  
19 to television commercials. In a preferred commercial mode a relationship is created  
20 between iamworthit and one of the providers of commercial insertion technology (at the  
21 head end). All cable providers not utilizing the particular commercial insertion  
22 technology provided by that particular partner provider would be subject to the use of the  
23 commercial replacement technology which iamworthit provides by delivering typically  
24 (but not exclusively) digital set top terminals streaming (HTML based) video  
25 commercials or the same preloaded onto the set top for subsequent display. The system  
26 is able to recognize if/when commercial breaks occur during the course of standard video  
27 programming through tags which are typically digitally encoded or (in the case of analog  
28 TV) transmitted within the vertical blanking interval.

29 The system could also allow a vendor to deliver (again by predefined rules) certain  
30 criteria for undercutting, in real time, any offer which the user browses (or even receives  
31 via e-mail). As in the case of time of purchase competition, these dynamic counteroffers

1 can be delivered to the user via either the advert replacement system or a pop-up  
2 window/java-script window.

3  
4 Vendors might even compete to "buy" the rights to exclusively deliver competitive offers  
5 to any offer which the user may happen to view and which is competitive to another  
6 product/service which the vendor provides. The terms of the rights which the vendor  
7 actually purchases may include the number of competitive offers which the vendor can  
8 provide in the future, the degree of the conditional "value" which the vendor would  
9 provide through the competing offer, the category(s)/domain(s) which the vendor  
10 purchases the right to compete, the degree of exclusivity/non-exclusivity. A vendor may  
11 also wish to constrain the ability of potential competitors to submit counter bids.

12  
13 Vendors in all other product categories can retain the ability to compete for any offer which  
14 the user receives. It is possible that vendors with exclusive rights may also try to lock a user  
15 into *future* purchases. Again, this ability to lock users into future purchase commitments  
16 may be auctioned to competitive vendors. The dedicated currency auction web site (e.g.,  
17 [www.creditauction.com](http://www.creditauction.com)) is architecturally very similar to the time-of-purchase competition  
18 variation of the service with the exception that it is a market place in which many users may  
19 submit their profiles with request for offers from vendors across any or all categories  
20 relevant to their profile as requested or agreed to by the user. A variation of this dedicated  
21 currency involves a scheme for time shifting into the future the transfer of funds from the  
22 users account to the vendor's (thus allowing the users money to accrue interest during that  
23 period). Thus dedicated currency whether it proceeds or preceeds actual purchase events  
24 guarantees the vendor sales which can be used -----plan capital improvements, attract  
25 credit or investment funding while providing benefits such as -----  
26 and/or interest on money spent to consumers.

27  
28 This reverse auctions for time-shifted purchases may also encourage vendors to form  
29 virtual retailer communities, to accept a common currency and offer a full range of  
30 products and services to a user. Any offers involving the user dedicating any form of  
31 "value" for the user for future purchases of products/services is applicable towards









accessible to the consumer, which reveals general and specific instances of miscarriages by health insurance carriers of their fiduciary obligations to provide timely payment for essential health care services. Based upon data regarding the insurance company's stated policy and statistical analysis of the ultimately approved or disapproved coverage, performance criteria may be determined pertaining to the percentage likelihood of the insurance company covering certain services and length of hospital stays for future services to patients. In order to provide a robust and comprehensive statistical analysis, it is also required that a certain core sample of non-coverage denied service be collected in order to determine a variety of correlation criteria which affect the probability of coverage in addition to the identity of the carrier, also the type of prescribed medical care for associated illness wherein coverage for hospital stay or treatment is denied. Other correlations are possible, e.g., the medical history of the patient, the family history, etc. Of additional importance, statistics may also be provided which indicate the degree of adherence (integrity) of an insurer to provide coverage and to provide it with the level of depth of coverage which was purported by the insurer according to its standard claims policy. The iamworthit insurance service is accordingly designed with two primary case objectives:

- 1) Provide detailed statistics based on robust historical data regarding realistic expectations which the patient can anticipate regarding acceptance, denial and quality on the basis of various types of coverage given the context of the insurance needs and requirements of the buyer and based upon the specific services and depth of coverage purported by the insurer and any associated medical information regarding the buyer. The service could even anticipate, by review of the collective medical statistics the most likely type of coverage the buyer may need, based on his/her medical record specifically within the context of the associated determined probability of coverage and quality thereof for each insurer in view of this medical history of the buyer. In addition, iamworthit may provide a general informational resource where buyers can access such statistical information by carrier. Data mining tools may enable the buyer to extrapolate the desired statistical correlations as to his/her type of profile of medical or specific features of his/her medical condition or specific type of coverage

1 and thus in light of this robust data better equipping the buyer to determine which  
2 insurance provider and associated policy is most appropriate for his/her specific  
3 health care coverage needs;

- 4 2) This information may be accessed in real time by the buyer (or for e.g. buyers,  
5 employer) while browsing specific insurers' sites, receiving specific offers by  
6 insurers (typically requiring certain medical information) or browsing an insurance  
7 retail portal. Typically, buyers of the present service will combine their use of the  
8 present browsing-based information with the Time-of-Purchase Vendor Competition  
9 Services (described below). Additionally, a rating feature may also be provided  
10 which enables the buyer to rate (and subsequent buyers to observe the averaged  
11 ratings) the various vendors by various types of criteria relevant to the buyer's  
12 experience in receiving acceptance, quality and depth of coverage (as well as how  
13 well the insurer adhered to promises for coverage as stated in its standard policy for  
14 coverage), of course, this feature could be used if for insurance products on-line for  
15 later purchase off-line.

16  
17 Iamworthit's time-of-purchase vendor competition application to medical insurance (as  
18 applies similarly to numerous other e-commerce commercial domains) involves the  
19 following protocol:

20 The iamworthit enabled consumer provides specific needs or requirements as to what types  
21 of medical insurance products or areas of coverage are desired. These requests may be  
22 submitted either to a portal (or "mall") which connects the buyer to an extensive resource of  
23 medical insurance providers, or, alternatively, these requests may be submitted to  
24 iamworthit which provides its own default portal to these resources (which likely exists and  
25 is marketed as an independent marketplace for insurance purchasers on its own behalf).  
26 Upon submitting these buyer requests, iamworthit is able to provide numerous competitive  
27 insurance providers offering similar insurance products which in this application may be  
28 accompanied by certain credentials (such as the buyer's age and many other clinical  
29 parameters).

30 Additionally, the pseudonymized digital medical records of the buyer (which may also be  
31 requested by an insurer, typically in addition or instead of submitting a questionnaire

1 regarding the buyer's health such as certain pre-existing medical conditions). Typically,  
2 patient medical records are a prerequisite to submitting an offer. The insurers may then  
3 provide specific offers for their respective products. Additionally, the buyer may browse  
4 various insurer's sites, providing pseudonymous medical data to the extent (and type) that it  
5 is requested and, specifically, to which providers the buyer desires. As offers from  
6 insurance providers are collected (in response to requests and/or browsing interactions), the  
7 time of purchase vendor competition is elicited whereby each insurance provider is  
8 automatically notified as to the most valuable offer(s) provided by a competitor(s) and are  
9 given the opportunity to submit counter offers accordingly. It should be noted that the  
10 determination of the value may be in accordance with certain buyer prescribed  
11 specifications such as price or certain expert determined criteria which consider either the  
12 price in combination with completeness and depth of coverage and, if relevant, the level of  
13 satisfaction of the buyer's request criteria.

14 Additionally, iamworthit provides the ability for buyers to share the information of offers  
15 they have received in accordance with the associated medical profile and/or response to  
16 health related questionnaire information, in order to determine if an offer which is submitted  
17 by an insurer inappropriately discriminates against the current buyer in terms of price or  
18 depth/comprehensiveness of coverage, or where such discrimination occurs in conjunction  
19 with buyer medical data which is prohibited by insurance regulatory guidelines as a valid  
20 criteria for such discrimination. Iamworthit's insurer coverage statistics (which the buyer  
21 may typically wish to observe in conjunction with the present time of purchase vendor  
22 competition functionality) are an additional important value determination factor (which  
23 may be factored in as an additional weighting criteria as numerically prescribed by the  
24 service (or by the buyer) and/or accordingly for purposes of setting a threshold by which  
25 insurers become disqualified as candidates for the buyer's business. In a variation, as is  
26 provided by the current model, buyers may instead first receive the historical offer  
27 information from other "similar" medically profiled buyers in order to submit a bid to that  
28 carrier which offered the most relevant and beneficial offer as well as all of the other  
29 competitive insurers. As is also provided according to the current scheme, the present  
30 iamworthit time of purchase vendor competition facilities may operate fully autonomously  
31 on behalf of the buyer (incorporating, if desired, specific request criteria).

1 Thus, optimal offers are provided to the buyer by a "buyer agent". Additionally, the present  
2 service may with some modification be provided to buyers dialing into an 800 number.  
3 Typically, in this case, however, (as is also the case in the on-line implementation), if the  
4 buyer does not have access to his/her pseudonymous medical records, he/she would  
5 typically answer a questionnaire for each provider (which would typically consist of an  
6 aggregate questionnaire satisfying the requests of all relevant carriers which could then be  
7 submitted automatically, or an agent could be programmed to automatically complete the  
8 questionnaire as needed.

9  
10 *Commercial Alliances with Regional or National Health Care Organizations –*

11 Because of the tremendous restrictions in providing insurance coverage to organizations  
12 which provide health care services, it will be greatly to their advantage for patients and  
13 physicians to be able to benefit from more fair and competitive insurance coverage by using  
14 time of purchase competition and the associated coverage probability service (providing full  
15 disclosure of non-coverage and inequities in delivery of coverage). The primary joint  
16 marketing entities which will benefit most significantly from this service and thus will be  
17 the strongest commercial allies, include independent physicians, medical clinics, large  
18 physician groups, pharmacies and perhaps even pharmaceutical companies, providers of  
19 medical supplies, hospitals and home health providers. Of the above entities listed,  
20 physician groups, clinics, and hospitals do a tremendous amount of advertising directly to  
21 consumers (and this is becoming increasingly true via the internet for physicians in private  
22 practice), particularly through the direct advertising medium of the internet (e.g., advertising  
23 on their web sites as well as email). These organizations, in light of their direct marketing  
24 exposure to the consumer which they provide, would likely become powerful allies in  
25 assisting the promotion and exposure of the competitive insurance with full disclosure of  
26 claim information service to their consumers.

27 Secondly, health care organizations which market to health care professionals, such as  
28 physicians and nurses, including pharmaceutical companies, pharmacies, and home health  
29 care providers may indirectly benefit by informing their target customers about the above  
30 consumer oriented service.

1    *Quality Ratings of Medical Suppliers*

2    iamworthit's insurance application provides a natural entry into a second commercial  
3    application of the service. In a fashion similar to that of Time-of-Purchase Vendor  
4    Competition Scheme, for purchasers of insurance products, another useful commercial  
5    application of the present service is facilitating the time of purchase vendor competition  
6    between medical suppliers and equipment providers. The synergy by which this second  
7    commercial application is enhanced is as follows: physicians and hospitals will likely be  
8    cooperative allies to the service in providing patient information substantially as required as  
9    well as promoting (along with home health care providers) particularly the iamworthit  
10    insurer coverage probability service to their consumers. They will also likely utilize such a  
11    resource for internal uses and if/when HMO service liability legislation becomes enacted,  
12    for legal purposes as well.

13    To the buyer's advantage, it is possible (as in the insurance application, albeit less likely)  
14    that unfair or arbitrary price discrimination may be detected and counteracted by  
15    iamworthit's data sharing feature. Common buyers in this environment include hospitals,  
16    physician clinics, home health care providers and to a lesser extent HMOs and end-buyers,  
17    the nature/types of supplies and equipment which tend to be purchased by each of the above  
18    categories of buyers are provided in further detail below. As suggested earlier in this  
19    description (and in issued patent "Pseudonymous Server for System for Customized  
20    Electronic Identification for Desirable Objects"), a vertical portal as the one described, may  
21    be personalized based upon the user profile as disclosed upon accessing that portal in  
22    general, at the item level as within a category (or in conjunction with search results).  
23    Additionally, vendors may be equipped with price discrimination tools and are in fact highly  
24    motivated to form a coalition by entrusting their data and counter offer responses to  
25    incoming time of purchase vendor competition offers and bids to SDI (or at least an industry  
26    specific counterpart of SDI) for medical suppliers in order to provide limits or "price  
27    ceilings" for any given buyer and associated profile and condition relating to competitive  
28    offers or bids.

29    As in the case of almost any category of purchaser, particularly the larger entities, such as,  
30    unions, hospital systems and large physician organizations, additional bargaining leverage  
31    can be applied to suppliers because of the volume of supplies which can be purchased. The





1 iamworthit medical supplies application as it similarly could apply to all of the above  
2 entities.

3 (e). Home Health Care Market - Equipment and supplies which are, in part, similar to large  
4 physicians clinics as well as some hospital - type supplies, used to provide out-patient care  
5 and maintenance services.

### 6 *Patient Medical Records*

7 The secure data interchange provides a secure and privacy protected storage and retrieval  
8 architecture in which buyers (in this case patients) are able to control the use of their  
9 personal profiles (in this case medical record information). Herein, we propose an  
10 alternative model which is based upon the assumption and belief that purveyors of this  
11 medical information, including hospitals, physician's offices/clinics, labor unions, (to some  
12 extent) HMOs are legally permitted to disclose and use their information for advertising  
13 purposes so long as a method is provided which insures that their information will not be  
14 associated with any identifying information about the patient (name, address, social security  
15 number, etc.). If regulatory constraints and/or the purveyor of this information prohibits the  
16 use of this information for the above purpose, patient permission could be acquired by  
17 patient signature on all appropriate consent forms at the time of admission and directly on  
18 the admission form.

19 In one variation, software may be provided to the organization which enables the  
20 organization to operate their own pseudonymous proxy server, thus enabling access by  
21 advertisers to the pseudonymized patient records. Email and/or telephony-based  
22 pseudonymous communications could be readily provided. For direct mail, the address  
23 information maintained by the hospital would mandate the hospital to provide direct mail  
24 advertising to the patients or to entrust this operation to a pseudonymous physical mail  
25 service such as SDI. In the (much) preferred variation, a trusted pseudonymous proxy  
26 server operator maintains the patient information on an external database (this would  
27 perhaps be SDI or an SDI-like service) in which all of the above pseudonymous  
28 communication media could be provided along with pseudonymous direct mail services.  
29 For the SDI data entrusted variation, the preferred business model to provide incentives  
30 to these organizations to cooperatively provide this data would involve SDI's privacy  
31 policy enforcement capabilities in combination with a revenue sharing model in which





1  
2 Essentially there is bi-directional information filtering: the requestor agent will only present  
3 certain information to the user, information that is relevant; and the requestee will only  
4 provide information when a request is judged to be legitimate. Information exchange  
5 between agents occurs as part of a multi-step negotiation, until both parties can agree on  
6 terms for either a physical meeting (or execution of a deal), or further pseudonymous  
7 exchange of information or cooperation.

8 In match-making applications the role of the central SDI query execution engine is to search  
9 continuously for "matches" between agents, based on agent profiles. When a match occurs  
10 additional information exchange may occur automatically between agents, until finally an  
11 introduction is made or an action executed by one (or both) of the agents. Decisions about  
12 what information to exchange are made on the basis of both static and dynamic profile  
13 attributes, e.g. standard (historic) profile information, current behavior, current location, and  
14 recent activity. Also mode of behavior e.g., social, business, leisure as detected by the  
15 user's calendar,, time, content in recent conversations, type of users recently interacted with,  
16 etc. LEIA style-behavior attributes can be used to automatically decide on the relevance of  
17 new virtual tag information. A requestee might also demand certain credentials to indicate  
18 the lack of negative reputation marks, for example that an interaction with the user has never  
19 received a bad rating (see Section 2.2.3). Perhaps a third-party could be used to determine  
20 whether the user's know each other (e.g. [www.sixdegrees.com](http://www.sixdegrees.com)). When a requestee denies a  
21 request for information, it may instead provide criteria for data releases. A requestor can  
22 respond with a different information request, or a subset of required credentials. Finally, the  
23 agents might agree on terms of negotiation and conditions can be anonymously fixed.

24 There are (at least 5 levels) of information disclosure: indicate to another user interest;  
25 release profile information; disclose communication; start a correspondence session;  
26 schedule a meeting/strike a deal, take some other action, etc. The end-result of information  
27 exchange could be an agreement to calendar a meeting for some future time and place; and  
28 absolute, or pseudonymous revelation of identity.

29 An initial implementation of the data-release policies might allow only manual definitions.  
30 However, after an initial "beta testing" phase, a data mining suite could be used to cluster  
31 users and generate exemplar data release and data request policies. A system can provide



1  
2 Some standard credentials which may be of interest to many users, and which may (as with  
3 resolution credentials) be incorporated with the standard settings of the user's data request  
4 policy as herein described. A few examples are cited (among countless potential others):  
5 profession, awards, honors, alma mater, e. g., Harvard graduate, doctorate degree, etc. In  
6 accordance with the parent issued patent US Patent #5754938, various credential issuers are  
7 provided for issuing standard and resolution credentials to individuals. Thus certain entities  
8 may be entrusted with "legitimate authority" to validate and submit credentials which are  
9 issued to the appropriate individuals. If a resolution credential is not issued (or not renewed)  
10 an adjudicating third party is provided which has access rights to both of the parties is  
11 provided to resolve resulting disputes (from the subject user). The present invention  
12 describes how credentials can be issued to users pseudonymously.

13 There are a variety of rules which a user's data disclosure policy and data request policy  
14 may contain, to control what if any attributes are released, and what credentials are  
15 required. A data request policy may state a rule for explicitly notifying the user if a  
16 particular resolution credential (e.g., indicative of a serious problem or concern) cannot  
17 be presented in response to the user's disclosure request. We allow initial information  
18 exchange to be anonymous, such that information that is released as preconditions for  
19 release of further information is not useful. Similarly, so long as initial encounters are  
20 anonymous there is no need to withhold information about them from the user.

21  
22 Some users may not wish to disclose specific information about themselves via these  
23 standard credentials but instead certain "extracted" more general information may be  
24 provided about themselves. For example, instead of a "Harvard grad or Ph.D." there may be,  
25 for example, credentials indicating "intellectual" or "prominent intellectual". Or instead of  
26 indicating an individual's wealth or value of assets, the credential may indicate "wealthy" or  
27 "very wealthy" (typically, depending upon user's wishes this latter credential should also be  
28 withheld during initial introductions or subject to some fairly stringent conditional criteria  
29 from the other party) and instead replaced with an even more general credential e.g.,  
30 "prominent" or "influential citizen"). Similarly, an individual's exact profession or scope of  
31 work may not be fully disclosed initially but rather a more general definition of his/her





1 behavior, and present temporal interest of a user and/or third party for purposes of  
2 employing the user's data disclosure and data request policies.

3 Credentials can allow users to identify other users that may pose a threat. This identification  
4 may be provided vis-a-vie resolution credentials and/or rating (by third parties). e.g. a user  
5 has not engaged in any serious criminal activity, physically harmed another person, or  
6 interacted with other individuals who are unable to produce these resolution credentials.  
7 Other credentials may specify the nature of an infringement, and its context and severity  
8 (e.g. what was the context of a physical assault? Was it performed during a bar brawl,  
9 against a friend, a boss, an elderly person, a child, a family member – or at work? In this  
10 case, the user agent may, for example, bring to the attention of a prospective employer that  
11 the user could not present a credential indicating that they had not previously harmed or  
12 threatened a former employer. Was it minor or severe? Also, if such individuals (lacking,  
13 for example, resolution credential proving the absence of having committed armed robbery)  
14 are (or come) within a certain proximity of a user, the user may wish to program his/her user  
15 agent to notify the user. The same would, of course, apply to a store clerk regarding  
16 customers of this sort or to baggage security personnel at an airport. Or, highway patrollers  
17 may be interested (e.g., on certain stretches of highway) in being made aware of vehicles  
18 and their locations whose agents are unable to provide a resolution credential proving the  
19 absence of a drug conviction.

20 In another application (in accordance with the auto insurance risk determination methods  
21 described in co-pending patent application entitled "Applications for Location Enhanced  
22 Information Architecture" [INSERT PATENT OFFICE NO]), an on-board computing  
23 device within a user's automobile could identify another automobile lacking, for example, a  
24 resolution credential for safe driving. i.e. the on-board user agent continuously polls agents  
25 in other cars for a "safe driving" credential, and if it fails to receive such a credential it  
26 issues a warning to the user. As an extension, this location data could be converted into a  
27 dynamic 2-D rendering upon the user's windshield (using heads up display technology) in  
28 order to thus superimpose a persistent flagging or highlighting of that particular automobile  
29 from the driver's visual perspective. Pedestrians or law enforcement officials (for example)  
30 could also receive instant notification. As is described in LEIA, a roaming cellular  
31 connection, or GPS, is not essential for providing a user identifier. For example, optically-





another potentially interesting individual for lunch, however, Mr. B's agent doesn't know all of the facts about how Mr. B may be willing to conform his schedule to meet Mr. C with regards to changing or rearranging a certain portion of his schedule, and/or replacing this portion of the schedule with another activity or engagement (an original activity which may now be in conflict due to the drive home and physical direction to visit Mr. C. Mr. B's agent thus identifies and conveys to Mr. B the salient changes to his schedule and the travel route within this possible alternative plan to meet Mr. C instead. Thus before Mr. A commits to Mr. B about lunch he has details about all of his feasible options for a lunch engagement. Mr. B may even state to his agent that that stop at the convenience store is a relatively fixed constraint (that it is important) and given that constraint what are the available options for a lunch meeting as such.

Note (as suggested) that at any time the system identifies within a certain probability threshold for that user that there may be a certain engagement (or other activity) which could be scheduled (or replace another schedule activity) the system:

1). Determines the predicted interest-level of that potential engagement or activity (or if it may involve replacing more than one activity/engagement, determine that the predicted interest of the combination thereof and/or it may enable other activities/engagements to be scheduled which may be replacements for others (estimate the predicted interest of the activities/engagements collectively comparing one scenario to the other). There may be other factors affecting user interest in one scenario versus the other, e.g. the distance or travel time with which it is associated. The system may also be uncertain regarding certain important determining variables, e. g. the profile of the user or even the present mood or activity or contents focus is not adequately determined, or the nature of the proposed activity or engagement is not sufficiently constructed (or robust) within the user's profile. In such cases, the user's agent can also ask the user direct questions in this regard using a natural language speech interface and (ideally) at the opportune moments in which the determination of certain facts as these within the user's present user profile would provide the ability to identify or dismiss for the user certain opportunities as they occur.



In another application of the system, the individual may wish to allow another user (or entity) to be able to arrive at certain levels of deductions or conclusions about the user which may concern them without explicitly revealing any facts or details used to arrive at those conclusions. Or in a variation, even allow a continuous persistent revelation of information regarding these certain conclusions such that if a certain conclusion(s) occurs (or ceases to occur) notify the requester of such facts. A statistics and/or manual rules-based approach could be used in this instance. A statistical interface provided to SDIs central data warehouse could enable an expert to establish correlations and confidence thresholds appropriate for deducing certain conclusions. For more novel or complex rules which haven't been seen before SDI could receive disclosure from users regarding their actions and a request as to desired types of conclusions. The industry could further trigger an anonymous action (and/or notification) on behalf of the user or depending on the privacy policy (consent of the other party) an action could even be triggered based upon such conclusion in the absence of disclosure of the conclusion and/or even the autonomous action taken.

E.G. notify the subject user of interest in message (or suggestion or warning) another relevant, noting third party, etc. In such cases, it may however be acceptable for the user to be able to receive fuzzed aggregate statistics regarding such data as a user's past behavior patterns (or more confidentially) other users who share similarities to the user regarding his/her user profile and/or perhaps context of similar activities and circumstances. User profile could be more general characteristics of the user than the detailed user profile as suggested herein, e.g., the user's corporate department peers, class, demographics, psychographics, travel venue or combination thereof. This process may be performed in either a pull or even push mode. Many associated applications are conceivable, for example, a spouse (the requester) may desire to have his/her agent persistently issued a resolution credential indicating that all of the activities of the user while on a business trip are consistent with behavior which is appropriate for a married individual, e.g., that the associations with business associates/colleagues are maintained at a professional level (per that mutual physical locations and perhaps even directly or transmitted spoken or written communications) or that user doesn't engage in late night carousing at questionable venues, or , for example, that the individual on his way to

work, to a meeting worked diligently on an important proposal or on his/her way home from work without revealing any further details which his/her agent has explicit knowledge of with which it deduces these facts with confidence. Another example might involve the use of LEIA via small cellular transmitters (with or without GPS) affixed to small children or the elderly whereby location and activity related information is maintained private, unless an inappropriate action is performed (e.g., motion and/or acoustic detectors in the elderly person's home suggest s/he may have fallen, that the child has wandered away from parents or is conversing with an inappropriate stranger. In such cases, notify the guardian or assistant in charge or allow a highly trusted party who happens to be in the immediate vicinity of the party to be notified to come to his/her aid. In the latter case (of the child) perhaps there is suspicion that the child has just disappeared and the parent(s) authorizing the agent of the child is able to send out an immediate description of the child to all certified "trustworthy" parties who are in the immediate vicinity, e.g., after the initial description has been sent out, one of the parties indicates the child was speaking or walking with an apparent stranger of X description which is also, in turn, transmitted to these trusted local parties (and at that point to the police). In a variation of the above example, a crime could have just been committed and the assailant has just been described in detail by the victims with information regarding location and direction last headed which is again transmitted to other parties: Thus, more immediate, complete and up to date information of importance for police to make a more immediate apprehension of the suspect is made available to police.

INSERT>>>>>>>>>>

Location determination via LEIA could be either through user's LEIA-enabled devices or the coordinates of the user(s) as identified verbally by the user, e.g., to a 911 server.

Other examples of the present system could apply to immediately locating trained medical, EMS trained or rescue personnel in proximity of a urgent medical emergency or accident. Such personnel could even be off-duty. Depending upon the particular nature of a medical emergency, particular specialists appropriate to the circumstances could be identified, their comparative physical proximities compared with other potential qualified individuals and based upon qualification appropriateness, physical distance and perhaps willingness, a most appropriate match made. In another application, likely witnesses to

1 an accident (including automobile accident), theft abduction based upon their location  
2 and time relation to the event of interest could be notified to provide helpful clues in  
3 assisting investigators in obtaining a suspect or legal professionals a conviction or  
4 acquittal. In the case of abductions timing in acquisition of such information could be  
5 extremely critical and the ability of such notification to be distributed instantly to all  
6 "trusted" individuals in the immediate vicinity at the very outset of when a person (e.g.,  
7 child) is even suspected of being lost could prove to be invaluable in not only rapid  
8 recovery but also in thwarting of those few instances in which a kidnapping has actually  
9 occurred. Some users may wish to be available and willingly notified to help in certain  
10 dire circumstances (the nature of which the user ultimately has control over) but may  
11 wish to remain anonymous to authorities and legal professional after the fact (or for  
12 example in cases of providing witness testimony (as via a subpoena) to more trivial or  
13 non-life threatening cases. A market model may also be created in which compensation  
14 schemes which are just appropriate to compel typical individuals to provide desired  
15 assistance may be provided. Finally, the present methodology could be readily extended  
16 to a taxi dispatching service in which a taxi's most local to a given user can be dispatched  
17 accordingly. Another example includes employees within a large organization. The  
18 statistical techniques as described in the parent patent application could be trained such  
19 that examples can be identified or explicitly stated rules provided by employers of what is  
20 considered to be actions and behaviors which are inappropriate and irrelevant to the type  
21 of work related activities of that employee's job description (ideally trained across  
22 previous "similar" employees). The system can take into account browsing,  
23 communications, (including spoken communications), even location (LEIA) data into the  
24 model (e.g., for tele-commuters) and flag suspicious behavior. SDI can even upon  
25 agreement by the employee) allow the employer to access more detailed information on  
26 the specific behavior once the suspicion threshold has been exceeded.

27 In another application, the user may wish to provide predefined rule(which could be  
28 suggested by the system initially in (similar fashions as suggested above).s which can be  
29 used to notify a certain individual(s) if certain explicit (or agent) inferred events occur,  
30 for example, that the user is in town, has just received certain important news of a  
31 promotion or change to a new position, internal company news which relates to that



individual and may have professional relevance on the other individual (the latter two examples which may be relevant to say the professional profile of that individual , etc.). In addition, the individual(s) may be made aware of the event disclosure policies set forth by the user relating to them specifically (if desired) and this individual(s) can further provide filtering rules which allow all or a portion of that information to actually reach the user. In another example, the user may have met or known the individual at some time previously and the venue and context of the meeting circumstances and relevant user profile information to that context are known and disclosed by both parties. The user then wishes to re-establish contact perhaps defining the reason for the request and the individual can have predefined rules or acceptance or rejection of the request directly. In a final example, virtual tags are provided containing comments and annotations regarding a certain physical object (which could be anything from a tourist attraction to the physical site of a recent hot news event to a useful annotation critiquing a bricks and mortar vendor). An individual whose profile suggests s/he is knowledgeable (or had provided the relevant annotation of interest) about that object would allow this relevant information to be conveyed to other individuals who have perhaps an interest in that information (generally or at that moment).

The relevant individual (if s/he is also willing) may be contacted and solicited to meet if their calendaring agents can appropriately coordinate the meeting. It may be either context specific, e.g., both individuals are physically local (or plan to be local to each other) and the physical object or site of interest or simply (generally) local to each other.

It is perhaps worthy to note that the present methodology could certainly be extended to include the application to rules dictating reachability conditions of users bu other users via standard telephone. The application to cellular telephony is of particular interest in its ability to communicate to the user in a variety of modalities (voice, notification, instant messaging, Web content). In addition, similar rules may be adapted in this scenario with perhaps minor modifications. For example, under what conditions may the user be reached by another user. E.g., what is the nature of the relationship between those individuals? If it is not known, inferences from the context of previous conversations may be determined. Does the user typically accept calls (or other communications) from the other user under

1 similar circumstances and/or times. If this data is not available are they accepted by other  
2 users similar to that user according to a key attribute(s) or does the other user accept  
3 communications from the user under similar circumstances as they present communication  
4 attempt (as this may suggest similar reciprocal acceptable reachability conditions as that of  
5 the other user. Again, LEIA suggesting location movements and associated measurable  
6 behaviors as sequential patterns and as formation of time may be important. This detailed  
7 statistical model of the user could be uiseful in learning the context of performing certain  
8 autonomous actions for the benefit and convenience of the user and in the reachability of the  
9 user by other users, e.g., in what communications mode is the user interested in  
10 communicating, e.g., speech to text-based instant messaging or standard telephony, does the  
11 user mind being interrupted with phone ringers, (or vibration mode) or ringer-off mode).  
12 Previous similar communications contexts may be useful in determining inferred rules  
13 which may be presently appropriate (e.g., was the spoken content of the previous  
14 communication indicative of acceptance or annoyance of the communication with the other  
15 user within a similar associated context . Eg., this may suggest whether and in what  
16 contextual conditions of both accessor and accessee the accessor's call may be patched  
17 through (or re-forwarded if initially accessing another number or telephone of another likely  
18 venue where the user may be believed to be present.

19  
20 Exemplifications of The Underlying Statistical and Rule-Based Intelligence Used in  
21 Autonomous and Semi-Autonomous Release of Personal Data Release, Matching and  
22 Reachability by Other Users  
23

24 The above section "Location Enhanced SDI System relating to Smart Home and Office  
25 techniques provides a relevant platform for integrating intelligence into a multi-device  
26 environment. Mobile user intelligence is a sub-component of this broader multi-device  
27 (or "ubiquitous computing") application-level intelligence platform. In view of the  
28 present application framework, it is a very intriguing and challenging problem to provide  
29 a statistical framework which is able to allow the user agent to make appropriate  
30 inferences regarding users who generally, though imprecisely, fit a certain profile which  
31 may suggest the activation of a rule (i.e., a "fuzzy rule") such as requesting further



constraints to the present system, e.g., precise accounting of present activities or focus (even such things as the likely emotions or present attention focus of the individual for which there may be a lack of sufficient measurable clues visa-vie the suggested input modalities. The natural language interface (perhaps an application for a language generation system) could communicate to the user its interpretation of his/her behavior and the statistical correlations with recommended autonomous actions. E.g., if the user is in his/her home study between 9-11PM, she/he is likely to be engaged in work relative to a research publication thus an invitation to other colleagues in this area will be extended and/or communication requests accepted to contact the user accordingly under these conditions. The system may detect similar conditions in which the user may reject requests to correspond by such individuals under these conditions, e.g., if the identity or identified context (or even context inferred from the identity) is inappropriate for that particular context of user activity. The user may provide another annotation to the rule which says that if she/he has recently corresponded with another individual(s) at that time/location, chances are that she/he is, in fact, in a social related mind set. As indicated, there are numerous potential inputs to the system, which could be considered in predicting what agent-mediated actions are, in fact, appropriate. The techniques of the parent patent application, if applied in this context, would allow us to make strong inferences about the particular mind set, reflecting present interests or preferences which the user is likely to be receptive to, presently based upon such clues as who the user is presently interacting with, the content profiles of the present real-time dialogues, e. g., typed or spoken through a communications network, (or simply passively collected off-line), the object profile of documents or web pages being interacted with. Activities may further be inferred by such indicators such as the particular physical location of the user, visa-vie LEIA, the particular sequence of location specific movement patterns, the particular types of devices with which the user is interacting )( and perhaps the nature and/or sequence of those interactions). More complex rules are certainly also conceivable, such, as the user will not interact, transact except under certain specific conditions, such as a commercial opportunity if the user fits a certain stereotype, i.e., similar to those previously interacted with for the nature or content of the proposed discussion, in light of the type of user, may be acceptable under certain circumstances,

1 locations or mind sets experienced by the user. E.g., a young female in her twenties may  
2 not wish to be approached by 40 year old males unless the context of the interaction were  
3 purely business, e.g., a professional, or commercial or sales opportunity thus if she were  
4 in a purely social mindset such requests to interact would be denied. Likewise a busy  
5 executive in his/her 30's may deny such requests from the young female unless her  
6 request was purely of a social nature and he was not overly time constrained. There may  
7 be certain "professionally" prospective correspondents, however, which he may prefer  
8 over others based upon some additional interesting criteria as well such as other  
9 business/professional areas of commonality, common interest areas, even socially  
10 compatible or interest criteria.

11 Of course, SDI is able to extrapolate extremely sensitive features regarding interest  
12 preference and context which may establish a basis where two individuals or more  
13 individuals, to be mutually identified to one another and/or communicate accordingly or  
14 to not be mutually identified and/or accordingly communicate (SDI may even establish  
15 the basis for these mutually beneficial interactions even if revelation of this basis by one  
16 or more parties by the other is not permitted within the privacy policy of one or both of  
17 the parties). Of course, SDI may frequently identify individuals whose personal data  
18 revelation requires certain minimal requested personal data from the other party as a  
19 condition to such revelation (which SDI is capable of fully automating on behalf of both  
20 parties) or the user may request certain information and make a personal judgment in real  
21 time as to whether further information about him/herself should be revealed and/or  
22 whether introduction or communication is desirable. In these cases it may be difficult to  
23 always successfully achieve an introduction when appropriate as the user is unable to  
24 guarantee

25  
26 a). The revelation of certain information to the other use if that user reveals certain  
27 information accordingly, or

28  
29 b). Is willing to agree to be introduced if certain desired information is revealed by the  
30 other party (often busy executives have absolutely no time to disrupt their busy schedules  
31 to discuss certain opportunities when there is enough revealed about those opportunities

1 to justify the correspondence. However, the nature of the opportunity is too sensitive a  
2 piece of information for the disclosing party to be able to reveal without certain  
3 guarantees that there is sufficient interest for the recipient to discuss the potential  
4 opportunity further. This quandry unquestionably poses a huge challenge, however, it  
5 can be somewhat overcome by:

6  
7 a). SDI's ability to keep track of most, if not all activities, behaviors and  
8 correspondences, particularly acceptance or rejection of those SDI pre-qualified  
9 recommended interactions,

10  
11 b). Detailed information provided by the individual or entity to SDI as to the nature of  
12 the various types of individuals, entities and potential opportunities to conduct business  
13 which the individual is able to provide to SDI. SDI itself may apply human experts for  
14 the task of interpreting the compatibility of certain requests with the access  
15 correspondence policy of the individual, provide a rules based system (which can also  
16 learn with fuzzy rules) with which to automate some of the decision or, at least automate  
17 the "definitive" exclusions).

### 18 19 **2.2.2 Applications**

20  
21 User and seller information - as detailed in the parent issued patent, specific details of  
22 what users and sellers may be looking to buy or sell respectively may be used to  
23 suggest the basis for a potential commercial transaction. The transaction may be  
24 large (but not necessarily so, e. g., real-estate, private investment in a small  
25 business or public stock). If a physical or on-line interaction with the other party  
26 is warranted (e. g., for larger commercial transactions), as is suggested later as in  
27 the present description, users may identify other users which form the most  
28 relevant "match" with their interest. At this point the agents can check for  
29 credentials, and then either communicate or calendar a meeting. Similarly, the  
30 agents may find the "best" match of users who happened to be physically  
31 proximal to the user at that particular time, or at some future time(s)/location(s)

1 which is mutually compatible (similar applications are suggested for matching  
2 sales persons with prospective clients, identifying experts to work (individually or  
3 collaboratively) on a particular project or problem, to answer a question of an  
4 appropriate specialized nature to their area of expert knowledge.) The parent  
5 issued patent suggests at a general level these commercial applications. An  
6 additional feature described therein involves the use of a decision tree called  
7 “Rapid profiling” which can be used in the present context to identify from the  
8 most common needs of users and “goods” of sellers in general and the known  
9 profile data about each user and seller individually, a list of questions for each  
10 party which most briefly and efficiently determines the complete user/seller  
11 profile of each party individually.

12  
13 Social Interests Profile Information—The parent issued patent also suggests the  
14 present application at a general level. For a dating application, users may be  
15 matched on the basis of their common interests/preferences and perhaps on the  
16 basis of certain information reflecting personality, social or cultural  
17 behavior/affinities or psychological attributes. On the other hand, for purposes of  
18 meeting casual acquaintances, users may be interested in another user who shares  
19 the above characteristics as well as someone who has recently shared similar  
20 experiences and/or personal challenges.

21  
22 Professional Information/Qualifications - As in the application of matching users and  
23 sellers, a description of a user’s needs or situation with relation to various  
24 professional services may be provided as additional data about the user.  
25 Examples may include: (as above) medical data, professional or business history  
26 (as well as legal history) which may be of interest to law firms, accounting firms  
27 or various business consultants. Personal, family or emotional difficulties may be  
28 of interest to psychologists or family counselors. Again, users may submit this  
29 information as a query for prospective matches, or they may be pseudonymous  
30 queries or automatically matched in accordance with criteria specified by the  
31 professional. The issued parent patent application also lists additional

1 applications, which could as well be relevant within the usage context of virtual  
2 tags.

3  
4 Employer/Employee Information - An employer may post a description as part of  
5 his/her virtual tag (and that associated with his/her company). His/her employees  
6 may also have provided ratings and/or annotations, which are further descriptive  
7 of his/her personality, leadership/management style and skills, work environment  
8 which s/he promotes and overall quality. A previous employee may allow  
9 him/herself to be contacted by the prospective candidate (e.g., in exchange for a  
10 fee).

11  
12 Access Privileges Information - Users in an organization are frequently given  
13 privileged access to certain files within a corporate intranet but not others.  
14 Though there are many ways of profiling users according to their level of access  
15 privileges to information, the following example is considered: Based upon the  
16 position (e.g., responsibilities and tenure with the organization), users may be  
17 "classified" into groups according to different levels of access to confidential  
18 information. Virtual tags may be used to extend the capability by providing for  
19 immediate disclosure of a user's information access privileges to another  
20 employee in real-time and in a physical context. Also, if a user reads or accesses  
21 certain information, meets with a certain colleague or friend, then a user might  
22 send a message X. This message could be (for e.g.) a request to perform some  
23 task relating to part of that information, a reminder to address certain issue(s)  
24 while chatting with the colleague etc. or, per the request of an individual's  
25 employer or colleague if a given individual (a sales person) meets with user X  
26 send him/her message Y (which may refer to a previous encounter, experience or  
27 fact s/he should know pertaining to user X and which may have bearing upon  
28 their conversation or professional interaction.

29  
30 Access restricted physical areas. In this case the virtual tag effectively may behave  
31 like an "electronic door key". A variation of the technique may be used for





benefit from this information as well as provide information for other vendor identification of competition. Other data resources including electric payment protocols, EDI, automatic check payment, check services, etc. may be useful data resources as well.

Again the disclosure of detailed business information is very helpful and a data release policy defining the parameters for such strategic initiatives may be critical in order to determine what companies may be potential candidates for which initial feelers (of high level information disclosure) would be appropriate to put out to a prospective company to determine mutual interest and/or further basis for expected synergies.

#### **2.2.4 Dynamic Annotation/Information Filtering**

In this extended application of SDI, we allow users and other third parties to annotate objects (physical and virtual) with meta-information, either to remind themselves about a previous interaction in the future—or as a system of “knowledge learning”, where systems of users leave useful information for other users. Information is left in the environment, leaving a trail for other users.

For example, the information that is tagged to an object, referred to as a “virtual tag”, can contain a pointer to other relevant information, such as a survey of a film by a third party, or the user’s own comments/feedback. For example, a restaurant listing could be annotated with meta-information about the quality of the food and service. Such information, when provided by a wide sample of users, can provide robust information about objects. The information that is used by a particular user can be filtered—for example, weighting the opinion of a respected restaurant critic, or weighting the opinion of users with common profiles (when that information is available).

Virtual tags (i.e. meta-information) can be assigned to objects with physical locations, and the information triggered based on the physical location of a user (using LEIA technology). Virtual tags can be assigned with expiration dates or other time-sensitive information. An individual user might leave an “action item”, for example—next time I return to this object (e.g. web page/ vendor) be sure to perform this task, enter this query, check this link for new information. As another example, after a conversation with an SDI-enabled user it is possible to tag that user with some notes, to remember the conversation the next time the two users meet.



1 audio pieces. The information itself can be embedded directly on physical objects, for  
2 example on magnetic strips or via. visual encoding techniques—or the appropriate  
3 information can be accessed from a secure remote database based on the user’s physical  
4 location (using LEIA location technology); or bar-codes that provide a universal identifier  
5 for an object.

6 As an extension to this model, we also allow users, vendors, and other third parties to  
7 associate “meta-information” with other *users and vendors*. This information might be a  
8 user’s opinion about his/her interaction with another user, an annotation that relates to a  
9 particular web page, or information about a physical object. The system of SDI enhances the  
10 value of this information by providing a secure environment where users can report meta-  
11 information (i.e. opinions) along with their profile information, to permit *collaborative*  
12 *filtering techniques* to generate appropriate meta-information about an object (user, physical  
13 object, vendor, web page, etc.) that will be useful to a particular user—given that user’s own  
14 profile. We define “virtual tags” as any piece of information about an object (physical or  
15 virtual). The information may be authored by any party, but annotated accordingly. For  
16 example, the appropriate virtual tag provided by a user about his/her-self is the  
17 pseudonymous profile for that user, -- and with SDI only the user his/her-self can gain  
18 access to the profile (either directly through editing, or indirectly through continuing  
19 transactions).

20 We might implement a Kasbah-style “reputation system” within such a virtual community.  
21 Initially users (under pseudonyms) have no reputation, and their opinion does not count for  
22 much, but after every positive interaction (as defined by other parties in an interaction), the  
23 “reputation” of a user can increase. (see the Kasbah system, MIT) [Kasbah 98] This  
24 reputation system is appropriate to a pseudonymous environment. Notice that gaining  
25 negative reputations is not useful when users can simply change identities. In one variation  
26 we can “block” certain users from providing information, when those users have negative  
27 reputations. Clearly, collaborative filtering or other data mining techniques could usefully  
28 allow for reputations when weighting information about an object.

## 29 30 **2.2.5 Meeting Planning**



A couple of concrete examples includes for example a real estate developer who develops clusters of home building sites or town homes which reflects the profiles of users sharing common or compatible profiles (e.g., socially, professionally or commercially) or recommending hotel lodging sites for individuals who share similar commonalties and happen to be visiting the same location or city. It is also possible to physically locate vendors at locations in which their most preferred customer prospects are most commonly physically traveling past or are physically situated (using LEIA) even considering where these prospects are during periods in which their shopper interest (or mood) for those items/services are heightened (or more generally factoring in both location and temporal (mood) factors into the user profile data as it is herein applied for this purpose. Whether users or vendors (subject to user consent) share customer data, there are other potentially interesting and relevant applications, for example, based upon LEIA a user's calendar schedule perhaps even verbal clues, it may be possible to coordinate meeting between users and professionals who can provide a useful service e.g., if a specialty physician happens to be situated in the user's present or anticipated locality and the physician would be available for an appointment at that time and location. It is possible that ad hoc use of shared clinic space may be available for such ad hoc appointments.

## **2.2.6 Investment Matchmaking/Venture Capital**

This section describes a market place in which start-ups can propose a venture or other financing needs in order to elicit an auction between different Venture Capital funds. We can allow potential investors to leverage securely confidential information regarding the details of business present and anticipated strategic alliances and customers. Each company securely registers all of this confidential business information SDI. It is even possible for confidential information about future R&D initiatives proprietary know how and intellectual property to be entrusted with SDI such that potential synergies may be determined well in advance of the market.

A domain expert(s) within SDI then determines potential synergy's between the various commercial entities both for purposes of facilitating introductions on a customer/vendor level, strategic partner level, as well as what particular financing sources which are compatible for the level of the financing needs and other characteristics of the financing









as limited facts in the abstract or possibly in more detail depending on the particular situation. There may also be particular critical junctures and impasses in the negotiation at which point the vendor may be compelled to concede. These points may be critical points at which it may be opportune for SDI to initiate or go into a deeper level of investigation with the other parties regarding the critical issues at hand. The above precedent may even be replicated among all other third parties, which SDI deems to be potentially receptive and legible to the prospective deal. This process is analogous to time of purchase competition in that if/when another potential deal is identified from another entity which is more desirable the original negotiating party may be informed (as well as possibly the other relevant parties). This process may be reiterated as well.

### 2.3 A Real Time Experts Market

In this section we suggest a market for expert opinions, in which users with an information-need are matched in real-time with “experts”, that are prepared to provide information and opinions in return for a payment. This economic approach can be used for example, if a vendor would like a user to provide feedback about its products and/or services: the user becomes an “expert” and can receive payment in terms of discounts in the future. The information in SDI can be used for the automated selection and user targeting of tasks, based upon profile information.

We can allow users to be identified by another individual in which there are mutually beneficial opportunities for both parties to interact and terms/conditions for the disclosure to the other party is defined within the user’s data disclosure policy. For example, issued patent 6,029,195 Herz, et al System for Customized Electronic Identification of Desirable Objects describes a “system methodology by which users are able to find a knowledge domain expert to answer a query, deliver personalized advice for a particular issue or problem to which they are extremely knowledgeable about, and provide references to other information sources.

The parent issued patent U.S Patent # 5754939 describes techniques for identifying experts on a communication network based upon their profiles and the search profile of a requester. In the present methodology we further provide an economic mechanism to encourage well qualified experts to provide options, and find experts. There is currently within the prior art economic incentive schemes which can be adapted and effectively

1 employed in this case in which it is possible to reward very well qualified (and truthful)  
2 opinions on a very broad range of issues and domains in which the accuracy of the  
3 ultimate outcome of a particular opinion is directly rewarded. This mechanism effectively  
4 insures the quality of the referrals/recommendations and is considerably more efficient  
5 then immediate payments for referrals where quality and truthfulness of the expert advice  
6 cannot be audited prior to payment, thus creating a disincentive to provide truly  
7 “valuable” advice.

8 The system of SDI can *forward requests for information to experts within the same system*  
9 *as time-of-purchase competition, SDI acts as a CLEARING HOUSE for requests for*  
10 *information, and experts can bid to provide responses.* The system of SDI can also help  
11 users to choose between experts, based on proficiency profiles that are derived from the  
12 value of information that experts have provided to other users. We can allow experts to bid  
13 for the right to annotate and provide ratings for particular pieces of information, products,  
14 services, etc.

15 In a web annotation example, as users provide annotations and recommend links, and other  
16 users provide feedback about annotations, then the system of SDI can build a “proficiency  
17 profile” for a user, to indicate the ability of a user to add value to the browsing experiences  
18 of other users. The information retrieval and document clustering methods as taught within  
19 the parent case #57549398 provide a statistically sound methodology to develop a user  
20 profile that predicts the “proficiency” of a user to provide recommendations about objects  
21 that fall in particular clusters.

22 In the economic variation we can allow agents to bid for the right to receive high quality  
23 recommendations, at least in the case where recommendations are new and cannot be  
24 duplication on the basis of current recommendations. We can create an “experts market”,  
25 where experts are assessed on their ability to provide quality to previous agents, as noted by  
26 those agents. SDI could also suggest a list of content categories that an expert might  
27 consider operating within. We can also apply the rapid profiling techniques in the parent  
28 case to assess a user’s expertise based on his/her response to a set of questions. A rapid  
29 profiling tree is essentially a decision tree, which can be used to present a sequence of items  
30 to which a user “responds”. The rapid profiling method enables a user profile to be  
31 constructed in the shortest possible sequence of inquiries to the user. In this case, the

response is the user's recommendation links (which are then judged). With the benefit of a comprehensive proficiency profile of the user the system may identify and present items to the user automatically (if s/he is amenable) for which the user is predicted to be proficient.

The market for experts and real-time information filtering can be applied to new content.

It can further be combined within the framework of the present incentive scheme to encourage users who tend to individually find news, which is of particular relevance to their daily lives and unique interests to proactively collect information (including but not limited to digital recorded audio/video) and provide that information to news distribution systems. Collaborative filtering leveraging both the overall quality/interest of user content as well as of the nature of the new content which that particular user is providing can provide a nice method to filter and channel this information.

The present clustering techniques may also be used to identify users that are able to provide useful new content. This content may range from commentary, opinions, critiques and





1 The criteria for matching the seller with a qualified prospective buyer is represented by  
2 client-provided meta-information associated with profile information, and can be used  
3 within a collaborative filtering system to determine the level of predicted interest which the  
4 buyer is likely to have for the items offered by a particular seller. The product space  
5 predicted to be useful to a buyer can be predicted based on his/her profile, browsing  
6 patterns, etc. Context can also play an important role, for example is the buyer currently  
7 engaged in browsing or buying related activities, can the buyer be persuaded to buy now  
8 with a good offer, or will the buyer never buy until he/she has performed more product  
9 search?

10 Vendors may also compete for sellers based on the profile information and success-profiles  
11 of a seller, and a perceived match with the products offered by the vendor. A seller may also  
12 work for numerous categories of products/services and vendors. Finally, the sellers in the  
13 market could enter into a market-based system, so that buyers bid for the right to work with  
14 a particular sales person.

15 Two other aspects of the present access to or by a prospective sales person include

16 (1) In an off-line context, we can use a LEIA-based method to identify buyers and sellers  
17 with similar profiles in physical space, and dynamically reroute their paths to allow a  
18 meeting to discuss a possible trade.

19 (2) Personal "Chemistry". In addition to a pure "product-space" set of profile features, the  
20 system might also consider wider compatibility between sellers and buyers in making  
21 meetings, for example choosing to introduce agents that share similar hobbies and spare-  
22 time activities.

23 The problem might also be informational: e.g. find an expert on ancient American  
24 civilization for purposes of writing an article, or answering a specific question. Relevant  
25 information might include the expert's resume, and the expert's knowledge expertise profile  
26 developed from his/her activities in responding to previous queries.

27 We might use a "fuzzy rule" to determine whether a user has a profile that is sufficiently  
28 close to an agent's goal profile to allow an agent-agent interaction and exchange of  
29 information. When the rule does not quite fire the system of SDI might also seek to clarify  
30 points of uncertainty, requesting further information until there is enough information to  
31 decide on the appropriateness of a contact.

1 It is clearly desirable to automate information exchange as much as possible, so that the first  
2 thing that a user knows about an agent-agent negotiation is after a deal has been struck that  
3 satisfies a user's preferences.

4 One approach is to perform data analysis on a large data set of users who share very similar  
5 profiles and to perform very similar actions and behaviors in all aspects of their successful  
6 interactions with other agents. The data analysis might also be used to suggest to agents  
7 when there is a problem, and when it will be useful for a user to provide more accurate  
8 information about what types of actions it is looking to perform in certain scenarios, i.e.  
9 when the current level of confidence about the action desired by the user is low. In many  
10 cases the system might be able to identify with a high statistical confidence level an  
11 appropriate action based on the collective statistics of other similar users and conditions.

12 When uncertainty exists the system can suggest the actions that seem reasonable, and allow  
13 the user to choose one action. A natural language interface might communicate to the user  
14 the system's interpretation of the user's current behavior, the statistical correlations with  
15 recommended autonomous actions.

16 For example, if a user is at home in his/her study between 9-11PM she/he is likely to be  
17 engaged in work relative to a research publication thus an invitation to other colleagues in  
18 this area will be extended and/or communication requests accepted to contact the user  
19 accordingly under these conditions. The system may detect other conditions in which the  
20 user may reject requests to correspond with these individuals, when the user's current  
21 context is inappropriate. Given this kind of feedback about actions, a user may now  
22 periodically review and adopt rules. For example, a user could also state that whenever a  
23 user has recently corresponded with a particular individual then the user is in a social mind-  
24 set, and would rather not engage in work-related activities.

25 There are numerous potential inputs to the system which could be considered in predicting  
26 what agent-mediated actions are, in fact, appropriate. The techniques of the parent patent  
27 application [INSERT US PATENT NO] allow us to make strong inferences about the  
28 particular mind set of a user; i.e. reflecting present interests or preferences which the user is  
29 likely to be receptive to, presently based upon such clues as who the user is presently  
30 interacting with, the content profiles of the present real-time dialogues, e. g, typed or spoken  
31 through a communications network, (or simply passively collected off-line), the object



profile of documents or web pages being interacted with. Activities may further be inferred by such indicators such as the particular physical location of the user (i.e. within a LEIA-based system [INSERT US PATENT NO]): the particular sequence of movements, the particular types of devices with which the user is interacting, etc.

In some cases SDI may identify individuals whose personal data revelation requires certain minimal requested personal data from the other party as a condition to an interaction, and SDI can again be used to automate this process if it is well explained in a user's profile. When confidence is low the user may request additional information (anonymously, via SDI) and make a personal judgment in real time as to whether to continue an interaction and reveal more information, and whether or not an introduction is desirable.

A simple but related application could be applied in the e-commerce domain to connect on-line customers with human sales representatives. A current version of this, offered by eStara, uses Internet-based telephony technology to enable browsers of on-line catalog pages to talk with knowledgeable sales representatives, simply by clicking on a “talk” button displayed on the web page of interest. Integrated into this methodology is the ability of the sales representative to identify which particular web page the visitor is

currently viewing, thus enabling him/her to more efficiently and intelligently assess the needs of the customer.

This methodology could be further enhanced through the tracking of the user's pseudonym and associated profile data provided via SDI, which represents a far richer set of information than that used by eStara (which only makes use of the current web page identifier). A customer's full profile can include demographics, previous purchases, previous web-site visits, physical measurements (for clothing purchases), health history, and income.

In a sense, individual sales representatives also have personal profiles; these can include experience, demographics, languages spoken, previous customer contacts, and product knowledge. An especially important factor is a representative's relative success in selling particular products – this history may reflect various personal strengths and weaknesses.

When the customer clicks on an on-line catalog's "talk" button, his profile is immediately forwarded to the sales system, along with the particular product he is viewing. Matching algorithms (as described elsewhere in the patent) then choose the most relevant sales representative currently available for on-line communications. The representative is presented with a screen containing facts pertinent to the sale (including product details and extracts of the customer's profile), and he is then linked to the customer via Internet telephony.

Extensions to this technology include:

### 1. Customized Price Discounts and Promotional Offer Recommendations

By combining the predicted affinity of a user toward a particular item with price elasticity metrics (gleaned from his/her transaction logs and from real-time feedback provided to the system by the sales representative (e.g., the customer appears to be in a

happy and spending mood)), it is possible to calculate a “personalized” price that optimizes the expected profit from a sale. This optimized sales price is communicated to the sales representative, who is authorized to offer the discount to the customer. Digital coupon technology would be used to ensure that the offer is redeemed by that customer exclusively, and within a pre-set time limit.

## 2. Automated Sales Representative Support

While interacting with the customer, the sales representative is presented with constantly-updated screen that provides information to support his sales effort.

### a. Rapid Profiling

If not much information is known about a customer, the representative would be guided through a script designed to elicit the maximum amount of information in the least number of questions. This information would be used to rapidly supplement the customer’s profile.

### b. Customized Sales Scripts

The representative’s scripted sales pitch would be dynamically adjusted to reflect the profile and current attitude of the customer. For example, if a customer is having trouble making a decision to buy (which could be detected by voice analysis tools and the total time elapsed in the transaction), especially persuasive text, geared to the customer’s demographic, would be presented to the representative.

### c. Additional Offers

If separate, but related, products are predicted to be of special interest to the customer, the sales representative is alerted so that he can promote the additional offers while still in contact with the customer.

In this variation, SDI can act as an exchange where professionals can exchange personal favors, for example one agent might indicate that it wishes to be introduced to another agent that can help with a particular professional problem. Both buyers and sellers can submit their respective needs, for example the need for a particular personal or professional favor, and in the case of sellers the ability to provide favors. Other personal information, for example: professional, political contacts, organizational affiliations, areas of competence, professional responsibility and spheres of influence etc. can be confidentially entrusted with SDI and used to make particular types of matches.

SDI can initially automatically define appropriate matches between the “buyers” and “sellers”. Given sufficient data, humans may use content analysis techniques to define rules based upon certain key examples or commonly occurring request and matching solution pairs. The techniques of the parent patent application US Patent # 5754938 can also be used to refine rules, and create new rules, for the cases in which the level of confidence with introductions based with the current rule set is too low. The mechanism for establishing the value at which favors are sold might involve a bidding scheme, but in this case professionals may bid a “personal favor” in return for another favor. I.e. the transaction becomes disintermediated from the outside economy, with a pure non-money based method for exchanging favors. A user with the ability to provide a favor receives bids from other users in need of the favor she/he is able to provide. Of course, the value of a given favor is likely to be worth more to some individuals than others.

The nature of the favor may range from very small and mundane, e.g., providing professional or personal advice or answers to a question, or it may be very significant and resource intensive from the standpoint of the provider, e.g., making a professional recommendation for a job position or political decision or providing recommendations/referrals for clients on a professional level or providing business recommendations to a partner or affiliate regarding a given individual or commercial entity with which to do business.

The techniques of peer-ratings within a reputation system can be introduced, and extended to cover a wide array of qualitative descriptive attributes to reveal the qualifications and ability of an individual, to allow some individuals in the bidding process to be excluded in

1 an initial “filtering process”, because they fall below expected standards of service and  
2 performance. An auction based model may establish a price for a particular type of favor,  
3 and for a particular individual, and therefore users with high performance ratings could be  
4 expected to sell for higher prices than other agents.

## 5 6 BARTER EXCHANGE FOR FAVORS 7

8 The qualification threshold of the provider of the favor is more important in some  
9 applications than in others. For example, if the favor involves an interaction with a  
10 particular individual for organization (e.g. a political favor) in which the individual’s  
11 anticipated ability and qualifications may be a significant factor in the likelihood to achieve  
12 the desired results. And thresholds may be set relating to quality or qualifications in their  
13 regard depending upon the relevance of this factor and/or the importance of the favor, which  
14 depends upon the type of favor requested. The bidding process is typically a reverse auction  
15 in which the request favor is passed around to multiple individuals (or entities) which SDI  
16 deems qualified for the particular favor requested (thus the lowest price which fits the  
17 qualifications criteria is typically the offer which is accepted. In other (perhaps many) of the  
18 barter exchange for favors, the qualifications of the requester may be an important factor for  
19 the requestee. These qualifications may include, not only professional abilities,  
20 proficiencies and credentials, but potentially much more subtle attributes relating to the  
21 individual's interpersonal, social or psychological profile and/or behavioral profile, for  
22 example, how SDI predicts that the individual receiving the favor will integrate and adapt  
23 within a particular professional context and/or interpersonal context which defines the  
24 situation and/or environment with which that individual must interact and/or perform, e.g.,  
25 as in a business deal, a professional opportunity, an educational opportunity, social  
26 opportunity, etc. In order for this pre-qualification of the requester to unilaterally work so  
27 as to assure the requestee with the level of confidence and trust in the requester which  
28 approaches that of a traditional favor (in which the requester and requestee are typically  
29 extremely knowledgeable and trusting in one another) requires a very efficient and perhaps  
30 fairly comprehensive reputation system revealing much about the above described personal  
31 and professional aspects of one of both parties. As suggested, in such a system and just as

1 important, is the trustworthiness of the individual providing his/her opinion about the  
2 credentialed profiles of the parties. Indeed a political or decision maker in a large  
3 organization would be extremely remiss (and perhaps even politically harmed) if an office  
4 or position were provided to another individual (requester) in which weaknesses in that  
5 individual's professional and/or personality later became evident to the detriment of the  
6 organization or political entity, as an example.

7  
8 It is even possible, that in some cases, the need may exist for a pricing model in which the  
9 individual providing these credentials through the reputation systems is compensated for the  
10 task. If such an individual must be fairly knowledgeable and assure to properly judge  
11 professional or subtle personality components, this individual may develop an independent  
12 reputation for effectively and truthfully judging these certain characteristics on a category or  
13 domain basis. If the subset of "judges" is fully distributed and ubiquitous permeating a  
14 substantial portion of the population as a default, if the qualifications of the judge are unable  
15 to be validated in advance, it may be preferable to provide payment following a reasonable  
16 level competency and success on the part of the requestee and requester. Or, in another  
17 variation, users are required whenever appreciated (and in accordance with their own  
18 judging qualifications) to provide opinions on the individuals (or other entities), and they  
19 must comply up to a certain quota and with truthful and accurate opinions in order to  
20 maintain privileges to engage in the service (e.g., [www.favor.com](http://www.favor.com)) or other applications or  
21 services resulting from the reputation system. It is also, of course, critical to ensure the  
22 judge possesses with regards to the individual or entity with which he/she is affiliated. In  
23 this case, the assumption is that the barter exchange for favors could cover nearly any type  
24 of favor desired by a user and the reputation system is implemented for most individuals.  
25 Of course, the present system may also be extended within a business context in which the  
26 favors include those provided by in between businesses. In both the individual and business  
27 application context, it is necessary for the above described use of privacy policies to dictate  
28 what types of individuals or entities a user or organization to include, exclude or price  
29 discriminations against e.g. to various degrees).

1 The following section "Agent-Mediated Value Exchange in the Supply Chain" talks about  
2 an economic system methodology in which a "value chain" is established for which entities  
3 may compensate other entities for present or often future value, which is likely to be  
4 sustained as a result of certain actions performed on behalf of that entity. This value often  
5 conveyed in the form of "barterables" could also be provided in the form of "favors"  
6 between the commercial entities (or potentially in a variation, individuals).

### 7 8 **2.3.2 A Market for Referrals**

9  
10 Situations frequently arise in a variety of contexts of human interaction (whether social or  
11 professional) in which a user may wish to refer the user they are in contact with to  
12 another individual. Often this occurs in a professional services context which a user has a  
13 particular need/or other characteristics which make him/her an appropriate match for the  
14 services provided by the other party. Or in a business context, often a user will forward a  
15 business contact or associate to another colleague who is deemed more appropriate for  
16 the particular context and/or scope of business. Likewise, in a personal or social context  
17 users may sometimes meet two or more individuals which they observe or perceive share  
18 common interests, goals or beliefs or perhaps possess complementary capabilities,  
19 knowledge, or characteristics. In each of the above scenarios, virtual tags may provide  
20 substantial benefits. For example, the referring user could forward the relevant portion of  
21 the profile and identified need of the user to the referring party whose user agent may  
22 determine the acceptability of the request and/or the priority with which a communication  
23 or meeting could be scheduled (e.g., as could be automatically arranged by/between the  
24 two party' calendaring agents). If the referring party's agent is unable to make a decision  
25 or priority assessment for scheduling purposes) on behalf of the user, the agent could  
26 instead try to contact the individual him/herself for assistance (and statistical feedback to  
27 the system's data model). In order for these types of referrals to be performed efficiently,  
28 the area of expertise required can be specified, and provisions can be made about the type  
29 of referrals that a professional will accept.









customized offers via coupons, through the use of cryptographically secure digital coupons which can authenticate the user and be delivered to the user in a non-transferable and non reusable manner.

Within the system of SDI we can also use secure user features in the SDI database to offer special offers to an individual. Of course, the vendor that requests coupon generation based on information about profiles can execute the request without knowing the profile or identity of the user.

As explained in the top-level SDI description, within the query-execution module rules , can be specified that determine whether a profile is suitable for an offer, and the type of offer to make; and the system of SDI can automatically furnish the provider of the profile information with a coupon (as long as that is compatible with that user's requests). The offer to a user can therefore be flexible, based upon a user's profile. The user may be given an encrypted code to present to a cashier when it makes a purchase, allowing offers to be made by vendors that are not on-line, to encourage a user to drive to their physical stores. More information can be available within SDI, for example to allow a user to receive information about quality ratings of vendors, for the products and services that have been provided to other users. In return, the vendor may request certain purchase pattern data from the user. In another variation, the user might enter certain identifying information such as the first several words of the offer, and then receive an email or fax of a barcode to take to the shop for scanning.

This method of discounting can be extended to users with Personal Digital Assistants (e.g. palm computing devices) in a store, that can use a scanner to enter bar codes of products, and then enter into an agent-mediated negotiation to secure a good price for the product. The user may be requested to disclose certain personal data, if a “better” offer is to be provided. Again, even in the in-store application the vendor does not need to receive actual information associated with the user, the processing can be performed remotely either at the user’s ISP proxy server or centrally on a SDI server, with the

1 vendor providing methods to adjust prices and offers based on a user's profile, but  
2 without seeing the actual profile.

3  
4 SDI may also implement time-of-purchase competition, soliciting and revealing to the  
5 user competitive offers from vendors in stores in a close physical proximity to the user's  
6 location. The user could also be prevented with other types of useful information,  
7 including: information for competitive products including (nearby) locations, price and  
8 functionality for eliciting time of purchase competition as well as previous customer  
9 complaints about each vendor.

10  
11 Finally, an alternative to dynamic negotiation with the vendor is that price labels in stores  
12 could encode a *strategy* that the store proposes to use to price-negotiation with a user, and  
13 just download this strategy to SDI and combine it with the user's profile to compute the  
14 final offer. The user receives a validated offer from the vendor, that certifies that the user  
15 with his/her current pseudonym is able to receive the price or discount. The offer can be  
16 encrypted to prevent fraud by the user.

17  
18 As suggested above, SDI acting on behalf of the user may selectively release only that  
19 information which based upon the presumed price/offer generation model used by the  
20 vendor will elicit a price decrease delivered to that user. Within this price negotiation  
21 process we can require that SDI releases just the right of information to optimize a user's  
22 price. This can be part of the understanding with a vendor (even though a vendor may  
23 commit to a strategy up-front). SDI cannot falsify information on behalf of a user, but is  
24 free to withhold information.

25  
26 It is certainly conceivable that such a technique could be deployed by vendors, to allow  
27 customized pricing for users as they shop in a store via information encoded with product  
28 identities, and negotiated via profiles stored in decentralized SDI nodes. A useful default  
29 to make the system work for non-SDI based users could be to allow a user to user  
30 averaged ratings on various criteria, annotations, etc.



system to track the average resale price, and determine an optimal time to provide the product for re-sale.

A variation of the resale market could include a “barter exchange” in which the seller receives another item, typically of comparable value, from another seller. I.e. agent 1 can state “I will barter A for B”, agent 2 can state “I will barter B for C”, etc. and the marketplace of SDI can look for a closed-loop of barterers, such that every agent is happy.

The user that buys a good can be an “advertiser” in a resale market, and other users subscribe and specify controls over the personal information made available to sellers/buyers. We can use a variation on credentials to act as “digital deeds”, such that the system of SDI validates when a product was purchased by the user, and whether it was new/used. This takes much of the uncertainty out of re-sale markets because then buyers are sure that the user is not falsifying information and items in the market. Digital deeds can act as credentials for ownership of a given product.

In fact, an agent can even be automatically asked by the system of SDI on completing any purchase whether it also wants to have the item listed and available for resale. If no, she/he is asked if/when at a later times she/he may change his/her mind. The system can request information from the agent about the price range that she/he would be interested to sell the item at, and the agent could be presented with a typical price range for the item. Of course, in a dynamic auction-based system this becomes the “reservation price”, i.e. the lowest price that an agent is willing to accept. Just as buyers can have more information about sellers, the sellers in the marketplace can have more information (even though it remains private) about buyers. For example, a seller can specify rules such as:

“I will only enter into a contract to sell to an agent that has never defaulted on payment”, etc.

Similarly, users (e.g. owners of kittens) may have certain personal interests for items of personal or sentimental value, and wish that they are sold to certain types of users. An interesting extension allows vendors to impose restrictions on the re-sale of items. The system of SDI can implement these systems for a vendor, so that at least within SDI products of certain types (e.g. software products) are not sold second-hand to other users for cheap duplication. Information goods can also be protected at source— for example with digital watermarks to prevent unidentifiable duplication.

1 Another variation is an auction scenario, where a seller brings an item to market, and buyers  
2 compete for the item. A professional, that makes a living by buying items and selling them  
3 on, might like an exclusive SDI-based auction site. Within the system of SDI we can support  
4 a *network of auction-sites*, that push agents around between sites and revenue-share, with  
5 source web-sites providing shares of revenue achieved in subsequent purchases on other  
6 sites. In another variation, the operator of a small site may be linked from a larger site, and  
7 provide a share of revenues to that larger site. SDI has a couple of important roles:  
8 monitoring purchases and ensuring that contracts are fulfilled, and also estimating the value  
9 of links if an up-front price is to be negotiated, based on information about the profiles of  
10 streams of customers to a particular site.

11 The information in SDI can be used to allow disintermediation of consumer-to-consumer  
12 markets, with agents able to reach ideal prospects based on personal information stored in  
13 the shared database, make an anonymous offer, and the consummate the trade of possible.  
14 This is likely the way of the future in this consumer-to-consumer market. Just as a vendor in  
15 current e-commerce systems can store information about its own user-base, and use that  
16 information to send personalized offers to users, within the system of shard information in  
17 SDI all agents can share information with all other agents, and the playing field is leveled.  
18 Users can query the pseudonymous user profile data and (if permitted), gain direct personal  
19 access to appropriate target customer prospect.

20 Conversely, prospective customers who are interested in a particular item(s) may query and  
21 access the database of users who are knowledgeable regarding that particular  
22 product/service (which may possibly be under a pseudonym), and also accessible to the user,  
23 i.e. available at an appropriate time and in an appropriate location.

## 24 2.6 Transportation Example

25

26 Another application of user profile database, in the case that the database contains also real-  
27 time information, is in an application to transport scheduling problems. The traffic systems  
28 in many large systems are congested. Via the system of SDI it would be possible to allow  
29 agents to represent the wishes of their uses within an SDI-enabled mechanism that  
30 controlled access to certain lanes on highways. The goal might be to control the flow of  
31 traffic, such that users agree to pool vehicles based on similarities across planned trips, and

1 users also payments in return for the right to travel at certain times and in certain lanes. The  
2 goal is a more efficient system that maximizes the sum value to all agents.

3 This is an example of a more general application area, where agents can look for  
4 opportunities to *change their actions and behavior in return for higher system efficiency,*  
5 *and/or payments.* We describe below an application to user agents negotiating on behalf of  
6 individuals ("travelers") to optimize the efficiency of a shared transportation system. The  
7 fundamental idea is to embed multi-user awareness into the basic fabric of a transportation  
8 system, in which the system aims to maximize the efficiency of the system by allowing  
9 payments between individuals.

10 To accomplish this task, the system must be able to anticipate both the immediate direct  
11 effects, a particular accommodative action will have on another agent(s) (mobile entities) as  
12 well as the indirect (cumulative) effect on the other agents and the associated time delay  
13 impact on each agent and on the entire transportation system as a whole.

14  
15 The techniques of the LEIA can be efficiently employed in tracking the present and  
16 anticipated activities, location and movement patterns of individuals who are in the process  
17 of traveling to a destination via simple or multiple modes of transportation, e.g., any  
18 combination of the following, i.e., pedestrians, automobiles, taxi, train, and bus (including  
19 public transportation). LEIA can be applied in both scheduling the flow of traffic for  
20 transportation media which are not subject to fixed time schedules (cabs, autos, pedestrians)  
21 as well as to dynamically improve overall efficiency of the movement patterns in order to  
22 assure that users on a collective basis arrive at their desired destinations with minimal delay.

23 The system of SDI can implement a real-time market where agents make and receive  
24 payments in return for changed actions from other users. For example, agents with high  
25 priority may be prepared to pay more than other agents for the right to enter a highway at  
26 any particular moment in time.

27 The market pressures at any point in time are between the marginal cost to an individual for  
28 changing his/her plan, and the cumulative marginal benefit to other agents in the system.  
29 Within incentive-compatible mechanisms, such as the regular price-based market place (e.g.  
30 the Walrasian tatonnement model) with enough agents, then it is optimal for an agent to



The central contribution of the SDI system is that it acts as a trusted auctioneer, receiving information from agents, and implementing an outcome based on that information without releasing that information to other agents. All that agents see are requested actions, and payments received or payments to make.

We could model a core sample of individuals on an active basis to determine implied valuation functions for different types of users, based on feedback provided by those users about how happy they are at any point in time. Within an incentive compatible system it is optimal for an agent to have a truthful and complete representation of its valuation function.

The problem is to discover that valuation function. Any help from SDI, via data mining techniques (i.e. similarities between my profile and the profile of another user that has rated his/her happiness) can be beneficial in reducing the costs of participation in such a scheme.

Then, before taking a journey, the user could provide some explicit information regarding the nature of the travel, for example the level of importance in minimizing travel time, whether the trip is business or sightseeing, the time of day, etc. With this information the system can assess the user's valuation function, and then have the user participate within the market, and secure a plan based on actions from all agents. The system might compute price thresholds, that represent different things the user can achieve for different prices. The system based upon its determination of situations/context can presents the price and time and urgency inferred for a user prior; which can be reviewed and corrected by the user.

With this profile information, the auctioneer can now compute prices to clear the market and maximize economic efficiency in the system; taking the following approach:

1. Armed with statistical information about the users of the transportation system e.g. the highway for the next hour, compute static (fixed) prices for different options. The goal is to maximize the efficiency in the market through a simple pricing mechanism, that is computed based on information provided to the auctioneer in advance.
2. Announce the prices, and allow users to use the system as they desire, with appropriate payments made. Payments can be readily extracted from agents via automatic toll systems.

1 Of course over time the performance of such a system could be optimized, as the auctioneer  
2 (the SDI marketplace) learns about consumption patterns of the agents over time. One way  
3 to provide incentives for agents to provide information to the system to allow up-front  
4 computation of useful prices to support an efficient outcome would be provide *discounts* to  
5 the agents that provide information. When such a system works well then roadway traffic  
6 can be controlled, congestion controlled, via an automated price-based system. Additional  
7 information made available from within SDI might allow adjustments to prices, based on  
8 unexpected flows of traffic. Just because the prices are optimized once up-front, based on  
9 projections, and announced to agents using the system; it can still be possible to adjust  
10 prices during any particular period—so long as there is a method for that feedback to alter  
11 the actions of agents, i.e. there must always be an alternative to make such a system work.  
12 Consider a two-road system, one shorter and one longer, then the longer route would be  
13 priced less than the shorter route, and the price differential adjusted based on current levels  
14 of congestion.

15 Again, to clarify, SDI has a key role in this system. Agents are only prepared to reveal  
16 valuation and trip data within a system that carefully protects their privacy, and controls that  
17 information.

18 Via the optimization performed to compute appropriate prices the model takes into account  
19 the effect on congestion that consumption decisions have, i.e. there is a relationship between  
20 the value to an agent that selects option A and the number of agents that select option A.  
21 The auctioneer needs a model of the transportation system itself to perform this type of  
22 computation. Essentially, we compute the fixed point of a non-linear system of equations.  
23 Stochastic optimization techniques are suitable for such a calculation.

24 As the system collects more data it can be less important to require explicit information  
25 from agents within the system, unless it is believed that there is a special situation about to  
26 occur (for example a Flyers game), which will have unusual effects on traffic patterns.

27 Co-pending patent application entitled LEIA-TR describes flexible tools and automatic  
28 traffic ticketing which could be integrated in to the same transaction infrastructure. SDI  
29 would provide a methodology to elicit information from agents, and dynamically set the  
30 price of traffic tickets and parking meter tolls to maximize the performance of a City's  
31 parking resources. For example, one goal of the transportation authority can be to optimize

the amount of revenues it can achieve, by charging more during “high-demand” travel periods.

Extensions of the present scheme could include coalition discounts, where coalitions of users can negotiate discounts based on group purchasing power (let us all use your system for a 10% discount or we will use an alternative road system). The coalitions may be formed automatically within SDI using shared profile information, allowing agents with non-competing interests to “pool their buying power”. Finally, it is interesting to note that users who share identical travel objectives can be encouraged to share vehicles, and/or public transport can be provided as alternative means of transportation on a dynamic basis.

### 3. Client-Side Data Mining Applications

Finally, we describe applications of *client-side data mining*, where the private information about an agent remains on its local client machine. Personalized of services and products is performed by pushing methods to the client machine, where the methods compute based on local profile information the most appropriate information for the user. This is very useful because the user retains absolute control over his/her profile information, but can still receive the benefits of personalization. This is an extension of iamworthit: users still provide profile information to the central database, to allow central queries and to receive value for that information, but never release identified information to a vendor.

#### 3.1 Client-Side Personalization

Client-side personalization allows vendors to push personalization rules to client machines that are special queries which use information store in a local client profile database to provide a user with personalized information and/or services. The personalization rules use locally stored profile information to compute an optimal product or service, or to allow a vendor to configure a virtual shop front. The answer to a query is returned to a vendor, to allow the vendor to push suitable commands to the client to enable appropriate displays to be created on the device with which the user interacts with the client.

Client side personalization is useful within the system of SDI because:

- a) It allows vendors to leverage profile information across multiple profiles that belong to an individual without explicitly receiving the information.
- b) It provides users with the convenience of using their true identities in payment and specification of delivery addresses for goods, because vendors do not need to





16 **One-time Client-side personalization.** In this variation the agent computer only  
ever performs a single user-personalization step, which is then maintained by the  
user as part of its profile for that vendor and passed in the future to allow  
personalization.

It can be computationally complex to compute the result of a personalization query. For example, profiles about the objects on a vendors web page or in its inventory can be as large as the object descriptions themselves, and full decision tree representations to decide how to target a consumer can be very large and complex. In cases where this is a problem, and it is inefficient or infeasible for the vendor to push a complete decision tree to the user we propose the following solution:

17 **Iterative Client-side Personalization.** In iterative client-side personalization the vendor and the user, via their agent computers, participate in an iterative distributed protocol to compute a personalization result for a user based on its profile. A similar method was earlier disclosed in US Patent #5753938. The idea is to structure the decision tree, and for example pass initially the first few levels to the agent. The agent computes the result of those levels, responds to the vendor, and the vendor passes the next few levels that are relevant to the result received. In this way, only the parts of the decision tree which are actually used to compute the result of a personalization query are exchanged with the consumer's agent computer.

Specifically for collaborative filtering applications, a simple technique can be used to reduce the amount of information that must be provided to a user to compute the result of a personalization query. In collaborative filtering a vendor's decision about which objects are









With a limited capacity to display an advert to a user, it is necessary to control the adverts that vendors display. We propose an auction-based system to compute prices dynamically through competition between vendors. The auction is a second-price sealed-bid auction that encourages advertisers to bid truthfully without “gaming” the outcome of the auction. The advertising auction is implemented on a user’s client machine.

In overview, the following steps are performed to support targeted advertising to users:

1. Buyer hits the web page of a vendor with an `iamworthit` advertising link.
2. Vendor pushes content to the buyer (possibly personalized).
3. Vendor contacts `iamworthit` ad server, requests an advert and pushes any profile information about the user.
4. The advertising server represents a number of companies that wish to advertise with appropriate users, and determines a short list of adverts to offer to the user's client based on profile information.
5. The user's client receives the choice of adverts and accepts an advert based on the values of bids and the likely suitability of an advert to a user.
6. The client requests the advert from the ad server, and then displays the ad.

Payment is collected for viewing the advert, and finally the user's response to the advert can be stored in the central SDI database for future tuning of advertising policies.

Figure 21 illustrates the system. The novelty provided by the system of SDI is that: (a) advertisers can decide when to place adverts and how much to bid based on historical information stored in the central SDI data warehouse; (b) the final advert is selected dynamically at the client with local profile information stored about a user. The stored profile information about users stored within Secure Data Interchange allows vendors and client-side SDI proxies to estimate the probability that a user will respond to an advertisement, given its profile and records about how other users with similar profiles have previously responded.

Vendors, represented within the advertising-network server have access to some information about a user, and can compute an expected value of placing an advert, based on the final purchases made by other users that were shown adverts. Meanwhile, client-side machines

1 have a more complete picture of a user's profile, and can compute the expected value of  
2 viewing an advert based on the advert and the price offered by vendors. In general, as the  
3 acceptability of an advert increases vendors will wish to bid more, and users will accept the  
4 advert for less money. Competition between vendors drives the bid price up, and allows  
5 users to receive the value of information about their profile to advertisers.

### 6 **3.2.1 Example: Technique to compute the expected value of an** 7 **advert.**

8  
9 A collaborative filtering system can be used to predict the probability that a user with profile  
10  $x$  will respond to an advert. The computation is based on historical information in the central  
11 SDI data warehouse for similar users to the advert, and also on the basis of historical  
12 information for similar users to similar advertisements if there is little information available  
13 about the actual advertisement. Information can be provided to advertisers anonymously in  
14 performing queries, and also randomized if necessary.

15 The basic technique is to select a set of users that have been shown the advert, or a similar  
16 advert, and are close in profile attributes to the current buyer. Call this the "relevant set".  
17 Given this, the probability  $\text{Hit}(x, \text{Ad})$  that a user with profile  $x$  will hit an advert with profile  
18  $\text{Ad}$  can be computed as the ratio  $\text{Hit}(x) = \text{Num\_Hit} / \text{Num}$ , where  $\text{Num\_Hit}$  is the number  
19 of users in the "relevant set" that responded to the advert, and  $\text{Num}$  is the total number of  
20 users in the relevant set.

21 Although subject to a certain level of necessary uncertainty, because buyer behavior cannot  
22 be predicted perfectly, ultimately it is useful to predict with as high an accuracy as possible  
23 whether a buyer will respond. Given a probability  $\text{Hit}(x, \text{Ad})$  that an individual buyer will  
24 respond to an advert, a vendor can define a bidding policy. The policy maps  $\text{Hit}(x, \text{Ad})$  to a  
25 value to bid for the right to target a buyer.

26 Assume in this section that the profiling system places buyers into one of several "classes"  
27 of buyers,  $C_1$  to  $C_n$ , given an advert with profile  $\text{Ad}$ , where each class has an associated hit  
28 rate, i.e.  $x \in C_1 \Rightarrow \text{Hit}(x, \text{Ad}) = \text{Hit}(C_1, \text{Ad})$ . Furthermore, assume that the system also  
29 predicts the average amount spent by a user that hits on the advert, or the average value to  
30 the vendor from a hit, this can again depend on the class of buyer, and can be determined  
31 within a profiling system based on historical information about this advert, or about similar

1 adverts. Let **Rev(C1, Ad)** denote the value to a vendor with advert **Ad** of getting a hit with  
2 a buyer in class C1. Now, the vendor can determine its expected value for placing an advert  
3 to a user in each class, for example using the computation **Val(C1, Ad) = Hit(C1, Ad) \***  
4 **Rev(C1, Ad)**.

5 The system of vendor-side advert competition works as follows. Once a user hits a web page  
6 with an ianworthit-targeted advert, a choice of adverts are pushed to the user's client from  
7 the ad server. Each advert is represented as a three-tuple:

8 ( link to location of advert, bidding function **Bid(hit, rev)** , profiling function  
9 **Profile(x, Ad)**, value function **Rev(x, Ad)** )  
10

11 The profiling function **Hit(x) = Profile(x, Ad)** is used to place a buyer into the appropriate  
12 class of buyers, and compute the hit probability for a buyer, **Hit(x) = Hit(x, Ad) = Hit(C,**  
13 **Ad)**, where **x** is the profile for the buyer, **C** is the class that the buyer's profile places it  
14 within, and **Hit(C, Ad)** is the probability that a buyer in the class will respond to the advert.  
15 A possible implementation of the profiling function is to encode it using prototypical cluster  
16 centers for a buyer in each of the set of buyer class types, and then assign the buyer with  
17 profile **x** to the cluster that is closest (in some well-defined metric) to the cluster type. The  
18 metric does not need to be linear in each of the dimensions of a user profile, and in particular  
19 would be expected to be place more weight on terms that are important to the hit rate in a  
20 particular class of buyers.

21 As described, this method has the following useful characteristics:

- 22 a) the hit probability for a buyer is computed at the client machine, using the profile  
23 that is stored at the client for the buyer. The client machine does not need to  
24 release the user's profile, and the **user's privacy is assured**.
  - 25 b) The information necessary to compute the hit probability can be encoded in space  
26 **linear in the number of clusters**, which is much more efficient that passing  
27 information about every relevant user profile to the client machine.
- 28

29 Without computational/communication restrictions one might pass historical information in  
30 the form of (Profile, Hit/Miss) pairs for users that have been shown the advert, or shown a  
31 similar advert. A nearest-neighbor algorithm could then be used to find the relevant set of

1 profiles for a new profile **x** and compute the expected **hit-rate** from the ratio of users with  
2 similar profiles that historically hit/missed the advert.

3 The cluster centers approximate this solution, representing the average hit-rate of “close”  
4 profiles so that the client machine can simply locate the closest cluster center and use that as  
5 a proxy for the probability that its user will hit the advert.

6 A similar technique can be used to compute the expected revenue from a buyer with profile  
7 **x** that hits an advert. In fact, this information can be computed using the same method,  
8 simply by associating an expected revenue with each cluster center.

### 9 **3.2.2 Client-side Advertising Reverse Auction**

10  
11 The client implements an auction for the right to show an advert to a user. The auction is a  
12 Vickrey auction between all the adverts that are passed by the ad server to the client. The  
13 Vickrey auction (Vickrey, 1961) is a second-price sealed-bid auction. Given bids **b1, b2, ...**  
14 **, bn** the auction sells the right to show an advert to a buyer to the vendor that bids the  
15 highest value for the value of the second-highest bid. The Vickrey auction is useful because  
16 it is truth-revealing. The optimal strategy of each vendor is to bid its true value for showing  
17 an advert to a user. For example, with a profit-margin of 20%, it is optimal for a vendor to  
18 submit **bid = 0.2 \* hit \* rev**. If accepted, it will pay at most **bid**, and it will pay only  
19 enough to shave the bid of its nearest competitor. The auction is sealed, so competitors do  
20 not ever see the information in failed bids. Cryptographic techniques can also be applied to  
21 ensure that the bids are not inflated by the auctioneer.

22 A user can also define an acceptance function in the auction, which represents her  
23 reservation price to view a particular advert. Assume that the reservation price depends on  
24 the hit rate **hit** of the advert, and represent the price as **Accept(hit)**.

25 The client-side advert auction runs as follows:

- 26 1. For every advert **Ad**, Compute the hit rate of advert **Ad**, given the profiling  
27 function **Profile(x, Ad)** and the buyer's profile **x**, and the expected revenue from  
28 showing the advert to the user, **rev**.
- 29 2. For every advert **Ad**, with hit rate **hit = Hit(x, Ad)**, compute the acceptance level  
30 **accept = Accept(Hit)** for the buyer and the bid for the vendor with the advert, **bid**

- 1       = Bid(**hit**, **rev**). If **accept** > **bid** then reject the bid for advert. If no adverts  
2       remain, jump to step 7 (in this case no advert is shown to the user).  
3       3. Given the bids that are not rejected by the client, sort them in order of increasing  
4       value, and insert a “buy” bid from the buyer equal to the value of its acceptance  
5       level for the advert with the bid of highest value.  
6       4. Accept the advert with the greatest bid, and charge the vendor the price of the  
7       second-highest bid.  
8       5. Fetch the graphics and URL for the advert, and display the advert to the user.  
9       6. The client monitors the actions of the buyer, and records (to be later transmitted to  
10      the iamworthit server) whether the buyer responds to the advert.

11  
12      The buy bid in Step 3 will be less than the bid for that advert, because the advert was not  
13      rejected in step 2. However, this buy bid might be greater than the value of the second-  
14      highest bid for an advert, and is required to make sure that the price paid by the vendor that  
15      wins the auction is greater than the buyer’s acceptance level. We make sure that the user  
16      cannot cheat by bidding just below the highest bid received by requiring that the user states  
17      his/her reservation value before the value of bids are revealed.

18      A more general system for dynamic customized advertising might allow a vendor to specify  
19      a bid for each type of buyer, and also a maximum budget, so that the vendor that places  
20      adverts can maintain control over its spending.

21      We might also allow users to specify in their advert acceptance policy how many adverts  
22      they are prepared to receive a day, so that they are not inundated with too many adverts,  
23      even if they receive financial compensation.

### 24      **3.2.3 Numerical Example**

25      The advert auction message from the ad server arrives with a choice of three adverts, **Ad1**,  
26      **Ad2**, and **Ad3**. Each advert is associated with a profiling function, a bidding function, and a  
27      value function. The client machine computes the hit rate for each advert, based on the  
28      profiling functions and its local profile for the buyer. Suppose **hit1** = 30%, **hit2** = 5% and  
29      **hit3** = 15%. The client machine also computes the expected revenue if the buyer hits an  
30      advert, **rev1** = \$2, **rev2** = \$8, **rev3** = \$3. The hit rate and revenue are used to compute bids

for each advert, using the bidding function. Suppose that the bids are **bid1** = \$0.50, **bid2** = \$0.70 and **bid3** = \$0.40.

Now, the client also computes the acceptance level for each advert, based on the hit rates predicted within the iamworthit system. Suppose **accept1** = \$0.30, **accept2** = \$1.00 and **accept3** = \$0.30. The bid for advert 2 is rejected because it is below the accept value. The bids for adverts 1 and 3 are accepted. Now, the auction is constructed with bids **bid1** and **bid3**, and **accept1** because that is the accept-value for the bid with the highest value that is not rejected. The auction takes the bids (0.50, 0.40, 0.30), and sells the right to advertise to the user to vendor 1 for \$0.40 (the value of the bid from vendor 3).

Finally, the client fetches the graphics and URL link information for the advert from vendor 1, and transfers payment from the vendor to the user for the right to show the advert.

### 3.2.4 Tuning an Advertising Strategy

Although the optimal strategy for a bidding agent in a single Vickrey auction is to reveal its true value for the right to show an advert to the user, the vendors are bidding over a number of different auctions, and might have a limited advertising budget.

#### *Selecting appropriate users.*

One good approach is to start with a reasonable policy and then adjust it dynamically, based on feedback received from adverts using techniques from reinforcement learning (see [RN 97] for an introduction). Off line simulation with methods such as Monte Carlo simulations can provide good initial policies, based on information about a simulated population of users, which could be provided anonymously by the system of Secure Data Interchange. Users can be provided with incentives to reveal information anonymously about their acceptance functions, to allow this type of off line modeling.

A vendor can tune an advertising policy with a random sample of users, simulating the auction that runs on a user's client machine. Metrics such as the average number of times that an advert is shown to a user of each type, and the average amount that the vendor pays to show the advert provide information to allow optimization. For example, the analysis might show that although a particular class of buyers are most likely to hit an advert, there is also a lot of competition to show adverts to buyers in that class, and the average cost to show an advert is high. In this case a vendor can conclude that it is more cost-effective to

1 advertise to buyers in cheaper but less relevant classes. The analysis can also be used to  
2 check that the average buyer is not setting an acceptance level higher than the vendor's bid,  
3 which again would indicate that the policy should target a different set of buyers.

#### 4 *Selecting appropriate adverts.*

5 The system that we have described can be extended to allow a vendor to select an advert to  
6 display to a user based on the type of profile of a user. For example, a vendor can pay for the  
7 right to be one of N vendors that compete in an auction for the right to advertise to users that  
8 hit the web page of a particular vendor. We can allow a vendor to not only tune its bid to the  
9 profile of a particular user that hits a web page, so that only users with a good fit with the  
10 service offered receive the web page, but we can also allow a vendor to tune the advert that  
11 it shows to a user.

12 We allow a vendor to submit a number of different adverts and bid functions via the  
13 advertising network server. The client-side auction can be expanded to allow multiple bids  
14 from each vendor, where each vendor can submit a number of bids, but only pays the second  
15 highest price bid from another vendor. Again, the auction is truth-revealing for a vendor.

### 16 **3.2.5 Automatic Advert-Replacement Systems**

17  
18 Although there are potential copyright violations in any system which alters the content of  
19 information published by a vendor before displaying that information to a user, we comment  
20 that the above described system can be applied even with vendors web pages that do not  
21 subscribe to the advertising network in SDI. Adverts can either be replaced, or added in  
22 separate windows. This enables any user that subscribes to SDI to receive personalized  
23 adverts and revenue from advertisers, irrespective of whether or not pages are within the  
24 SDI network.

25 The key problem in ad replacement is to identify an advertisement within a page, i.e.  
26 information that is not related to the core purpose of the page. At present adverts may take  
27 one of two forms:

- 28 (a) They can be statically, or dynamically, generated at the web server of the web  
29 page that the user hits, and pushed to the user directly at the same time as the  
30 general content.



(b) They can be pushed to the user in parallel with the other non-advertising data on a web page, for example with a link to an advertising network.

The adverts in (b) can be identified by tracking ad server URLs that are embedded in the source code of web pages. Case (a) is more difficult because there is no identifying URL to indicate the content of the section of a web page. However, there are solutions to advert identification in these cases, for example via regular expression identification and the location of ad banners, as described in [www.junkbusters.com](http://www.junkbusters.com) and the “block files” that allow a proxy server at junk busters to strip adverts from pages. Another technique is to identify URL’s that allow a user to click to another domain from within the current page. These might likely be adverts, consider for example, a link to [www.buyanewcar.com](http://www.buyanewcar.com) on the New York Times web page.

Another variation applies to television commercials and other broadcast medium, especially as the data storage capacity of local set top boxes increases. For example, a code is transmitted by national networks to cue local adverts from local network stations, allowing the replacement by personalized adverts from advertisers in the SDI ad network. In the future it is also likely that we will use digital storage devices to receive entertainment and television programs that are broadcast, storing the information for future viewing, or even slightly delayed “on demand” viewing. With digital media and storage, it is possible to replace adverts inserted into programs by the producers with SDI targeted adverts, either with or without the cooperation of the source of the information.

### 3. Additional Applications

1. Use of Real Time and Anticipated Vehicle Location Data to Provide Real Time Traffic Reporting and Predicted Traffic/Congestion Modeling

a. Real Time Traffic Reporting

A simple but useful application of LEIA involves on an opt-in basis real time transmission of a sample of driver’s location information to a regional traffic-reporting bureau. Though existing prior art approaches are reasonably efficient, there are common

1 circumstances in which congestion may have just occurred (e.g., following an accident)  
2 or wherein congestion has recently cleared up and such up to the minute information is  
3 not available to the traffic bureau. Moreover, this approach could be extremely  
4 advantageous in providing optimal navigational assistance for drivers wishing to  
5 optimally reduce driving time to a particular destination, which she/he submits to the  
6 system. The present system would consider the comparative length of different routes  
7 from the user's present location to the destination in combination with the anticipated  
8 average speed based on that of other vehicles traveling those particular routes. This  
9 enables real time recommendations to be presented to the user, based on this real time  
10 data.

11  
12 b. Predictive Modeling of Traffic Patterns Based upon Real Time Location Data from  
13 Vehicles

14  
15 It is possible to apply certain statistical techniques in order to predictively anticipate  
16 traffic patterns for some marginal period of time into the future based upon real time  
17 location data regarding the present case sample of vehicles. One may consider recent  
18 past traffic patterns compared to present traffic patterns (i.e., change in speed as a  
19 function of time by segment of highway as well as absolute and rate of change in traffic  
20 volume). There are pre-existing models which are able to anticipate the corresponding  
21 effects of congestion at the moment in which it begins to occur which utilize these  
22 variables. This model must also consider the rate of change in average speed over each  
23 segment as well as the predicted impact of changes in traffic volume from feeder  
24 highways based upon changes in volume occurring dynamically thereupon. It is a key  
25 objective of the navigational assistance feature to utilize future predictive traffic models,  
26 in order to optimize the reduction in transit time for the most number of users possible,  
27 i.e., considering the volume of traffic which is predicted on a segment by segment basis  
28 in combination with the traffic models for the impact of these volumes upon transit time  
29 (including the rate of increase thereof). With this data, the system can determine, in  
30 advance, on an individual driver basis, which route for a given requested destination, will  
31 provide the shortest drive time for the user, while insuring that a near optimal transit

time is achieved, across all segments of the highway system, as a result of each individual system recommendation to each driver. For example, an anticipated compounding congestion problem would warrant the vast majority of vehicles to take alternate routes until the congestion or predicted congestion is alleviated in that segment.

## 2. Use of LEIA for Creating Traffic Models for Purposes of Mapping Commercial Industrial and Residential Real Estate Market Opportunities –

A very useful application of LEIA involves the collection of traffic pattern data on a time-specific basis (the day of the week and time of day) for the traffic, as it passes each piece of real estate. An electronic map which is ideally Web based (potentially nationwide is generated and constantly updated based upon this data. Additional information may be provided which may include (but is not limited to):

a. Origin and destination information of the vehicular traffic (as captured by LEIA) which may, especially if correlated with time, suggest the nature and context of the driver's activities, e.g., rush hour traffic, errand traffic, etc. It may be useful to factor in the type of neighborhood the vehicle returns to every night, the type of commercial or business entity she/he drives to work to each day, etc.

b. Other activity-related clues which the user is willing to release, e.g., devices interacted with, content interacted with or transmitted information, etc., which may provide insights into the mind-set of which users tend to experience when in the vicinity of the real estate property.

c. User Profile Data – Aggregate purchase and content affinities as well as price elasticity data (gleaned from purchase statistics) could be very useful information to commercial real estate developers and purveyors. Users with the right profile and a receptive mindset are of particular interest.

1 It should be noted that the present system may be extended to residential real estate. E.g.,  
2 what types of jobs (such as quality of jobs) do local commuters have? What are their  
3 numbers? How far do they commute (particularly if they tend to commute further than  
4 the present real estate site)? Do their commuting routes tend to pass the current potential  
5 real estate site? The last three questions would also be particularly relevant as well for a  
6 prospective industrial real estate development opportunity.

7  
8 Industrial real estate developers also may be interested in mapped models of real estate  
9 depicting the professional and known likely educational characteristics of the associated  
10 local residents in that region? What are the other businesses at which they work? (If  
11 available) what are their particular positions/responsibilities?

12  
13 3. Use of LEIA for Creating Traffic and User Profile Models of Traffic Passing  
14 Billboard Sites and Providing a Map of Such Information on an Available Billboard Site  
15 Basis –

16  
17 The presently described techniques for providing dynamically updated informational  
18 maps containing detailed statistical data regarding vehicular traffic passing real estate  
19 sites can be further extended to similar maps of interest to advertisers which contain  
20 locations for available highway billboards. The present system further provides Web-  
21 based access, which enables advertisers to make reservations and purchases of such  
22 billboards. In one preferred variation, an economic model is deployed to optimally price  
23 the billboards. I.e., a varied representative sample of each type of billboard sharing  
24 similar traffic/user profile features with others is auctioned for this purpose. In another  
25 variation this on-line auction model is deployed for all billboards available by the system  
26 in order to provide a novel service to advertisers which is a “billboard auction” site.

27  
28 4. Use of LEIA for Enabling Drivers to Identify and View Sales Opportunities and  
29 Offers Associated with Physical Objects in the Surrounding Environment Utilizing  
30 Heads-Up Display Technology

The present system is an extension of the virtual tag methodology in which physical objects contain meta-data and in which the location of such objects with respect to the user is determinable either by transmission of a signal, which contains the exact physical coordinates of the object or a scheme by which this information may be pre-loaded, which is typically from a server, which contains such information (thus “non-wired” objects may be identified accordingly) to the user’s device, e.g., as she/he physically travels in the vicinity of such items. In either case, the present physical location of the object must be known in addition to that of the user (in the case of fixed objects, of course, the dynamic location tracking techniques are unnecessary). A primary commercial application of the present scheme applies to purchasable items. The user may submit his/her interest profile to SDI for purposes of being dynamically notified and disclosed of purchasables which match his/her user profile or specifically requested search criteria. The metadata may include a brief description and any additional level of information which the seller may be willing to disclose to that particular user. The seller may possibly request details regarding price elasticity of the buyer (purchase appetite) before quitting any discounts to the existing price. Prime example applications of the present scheme may include:

- 1...Revealing used car sales opportunities to drivers (by actually flagging relevant vehicles which the driver passes (or which pass the driver);
2. Real estate and home sales opportunities (as well as apartment and commercial real estate rental/leasing opportunities;
3. User-user introductions in which one or both users may be extremely interested in the other based upon matching criteria within their associated social, information and/or professional desires and assets.

The preferred underlying technology for the present heads-up display system involves a technique for tracking the orientation of the user's head as well as the direction of his/her pupils (it may be possible, however, to implement the system by detecting pupillary

direction based solely upon the position of the pupils relative to one another, using a 3-D optical tracking device. The iris of the user may also be utilized as a bio-metric identity of the particular user (see issued patents assigned to IrisScan Corp.) In the preferred implementation, the head-up projection device projects the virtual flagger upon the portion of the windshield or window(s) which is in direct line of site between the user's pupils and the physical object of interest. The technique may utilize heads-up windshield display technology and/or remote projection (e.g., for projection upon the side windows).

In another more advanced version of the system, the holographic projection system using micro-mechanical techniques may project the information directly upon the user's retina. This technique is detailed in co-pending patent application entitled "Remote Retinal Imaging Projection System" The variety of other potential applications of the present technologies are also conceivable which include:

# SYSTEM FOR PROVIDING REMOTE ACCESS AND MODIFICATION CAPABILITY OF DOCUMENTS AND E-MAIL BY A SENDER

The present system involves the use of a technique which enables the sender of a document, upon permission of the recipient, to provide subsequent updates and revisions to that document remotely and automatically in the absence of the recipient. Utilizing the computer's modem and e-mail delivery system, these updates could be provided remotely by the original sender (who alone can gain appropriate access privileges through a unique access code or authentication). The sender may also be able to identify whether or not the old version had been accessed by the recipient and possibly even at the level of the particular segment, or segments in which the specific changes had been made. (This capability could also be two way in which certification of not only receipt, but also access particular portions of the information by the intended recipient could be positively confirmed by the sender however, subject to approval by the recipient). If the recipient is amenable, such receipt confirmation could also be automatically time stamped by a secure trusted agent on the recipient's machine. The system could also perform the appropriate modifications in different drives on which the document is stored. If the

1 appropriate modification was performed on the hard drive, but not a copy on the  
2 floppy(s), the recipient could be notified and prompted to insert the floppy(s) for  
3 appropriate updating until all versions of the original document are updated. It should be  
4 noted that such latter functionality would even be of independent utility to standard  
5 modifications as performed with word processing systems. The above same features  
6 could also be provided for e-mail as well(in a variation, time stamped receipt could also  
7 be applied to voice mail as well). The recipient if he/she had not approved the remote  
8 modification in advance, the sender could request the access/modification privilege of the  
9 recipient via e-mail or instant messaging service for example. The recipient may also be  
10 able to identify a digital time stamp for the relevant portion(s) of the document (or e-  
11 mail) as modified/updated by the sender. In a variation, XML metadata containing  
12 digital time stamps could also be utilized for the standard modifications, which occur  
13 throughout the document, i.e., the nature of the modification, i.e., the addition, change or  
14 deletion, the time it was performed, on which machine, under which password and if  
15 relevant by which remote user. Outside of the context of the present remote modification  
16 system this function could be integrated into a standard word processing system with  
17 standard comparison check function. Such digital time stamps would help in proving  
18 identity and date of authorship at a detailed level.

19  
20 In a variation of the present system, a user could also receive automatic notification  
21 if/when a recipient has accessed (or provided an associated modification to) a document,  
22 portion thereof, e-mail or voice mail. This would require the recipient to provide  
23 approval of the above privileges for the sender in response to a request prompt. In  
24 another variation of the present system, the provider of a URL may receive time-stamped  
25 verification of a user having accessed a URL (which may request a uniquely identifying  
26 access code or other authentication or the Web page may be accessed by that user who is  
27 provided with a completely unique URL, thus verifying access by that user for remote  
28 access, modification and associated time stamping. If the URL is not under control of the  
29 requestor of the verification the accessor's client or Web-centric script (associated with  
30 his/her access account) could notify him accordingly with similar time-stamped  
31 verification of access. In another related application of the present capability, the

accessor's script could carry-out certified verification of more diverse or complex actions by a desired accessor, for example:

1. Verify (and if desired notify) the requester of the accessor having sent or received and read document (or e-mail) with content X to or from recipient or sender Y. For example, knowing that recipient had read a message or document sent by the sender (or another sender) or otherwise accessed may prompt requestor to take the next action or communicate with recipient.

Or verify (and if desired notify) the requestor (e.g., an employer) of the accessor (e.g., an employee) having made telephone contact with X individual(s) or phone number(s) at a certain time(s) (or within a certain time frame and/or containing certain message-type or content). Or the content (contextual profile) of the conversation(s) may be revealed or confirmed as being within a target contextual domain(s).

2. Verify (and if desired, notify) the requestor (e.g., an employer) of the accessor (e.g., an employee) having performed certain definable and/or desirable volume of on-line tasks, e.g., typed certain type(s) of content, performed certain accounting, billing (telephone handling courteous) prompt or other administrative functions. Knowledge by requestor (e.g., via notification) may then prompt requestor to take a subsequent action (or e.g., communicate again with accessor).

3. For any of the above notify requestor (e.g., employer) of the failure of the accessor to perform certain desired (and/or volume) of tasks with pre-defined threshold parameters.

4. Verify and, if desired, notify requestor (e.g., teacher or parent) of accessor (e.g., student or child respectively) of certain on-line content consumed, e.g., Web pages, paragraph by paragraph content (including wireless e-books) or on-line homework, such as quizzes, essays, reports, in conjunction with or independent from such on-line content.





1

2 As described, are some primary examples of how the present technique can usefully be  
3 deployed. User profiling may also be usefully applied for purposes of collecting feedback  
4 about various types of user reactions and responses to various on-line stimuli (e.g. general  
5 Web and e-commerce sites to various products such as movies, music, interactive content,  
6 advertising news, interactions with other users etc. A somewhat related methodology was  
7 disclosed in the parent patent involving a technique for allowing users to rate vendors  
8 according to a variety of criteria and enabling the vendors (e.g. for marketing purposes)  
9 and/or users to observe statistical correlations between the user profile characteristics and  
10 ratings (using data mining techniques). In a simple example, application of the present  
11 methodology, the users viewing a Web page (containing any of the above suggested  
12 informational types) are allowed to provide their personal comments about the page or about  
13 their own experience in viewing the page. Informational retrieval and statistical NLP  
14 techniques may be used to cluster (using standard clustering techniques) both the comments  
15 (by similarity of their content profiles) as well as the users (by similarity of their user  
16 profiles). In the first instance, the comment cluster exemplars are extracted, used as the  
17 comments which most closely exemplify each of their associated clusters and thus presented  
18 to vendors and/or users. The aggregate (average) of the user profiles associated with that  
19 cluster, the user profile of the user who provided the exemplar comment or the subset of  
20 user profiles which characterize subclusters of the comment clusters (wherein the  
21 subclustering routine is based upon the profiles of users within each comment cluster) are  
22 also presented to the users or associated vendors in conjunction with the exemplary  
23 comment most closely associated with that stereotypical profile of users. In a variation,  
24 which could be an automatically selected alternative if the exemplar fails to provide a clear  
25 representation of the various comments in the cluster (if further subclustering does not  
26 adequately achieve this objective), it may be possible to "combine" these various comments  
27 which are closest to the clusters' centroid by identifying those portions of the comments  
28 which are similar in meaning and those portions which are different. For similar portions,  
29 by preferentially utilizing the counterparts which are closest to the centroid and adding to it  
30 those portions which are different in piecemeal fashion, the desired objective may be  
31 achieved.

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30

In a variation, the different types of comments may be clustered and accessible to users via a hierarchical cluster tree used to create a menu of automatically labelled clusters (see parent patent detailing this general method). Users may access such menus by combining querying (of user or content attribute with menu navigation).

It is of course reasonable to combine the technique in the parent patent application for statistically correlating user profile attributes with ratings associated with the users possessing those attributes.

It is also reasonable to apply the present technique to video (or streaming video content) whereby during the course of viewing the content relevant survey questions are presented to users and if desired, the correlation statistics of the results with the associated user attributes are presented to the content owner and/or (preferably) also future viewers. In a novel and advanced variation, time shifted video technology would enable the ability to capture audio or full video/audio versions of the comments of the user.

In a very novel application, it is possible upon permission of the associated users, to automatically construct menu trees (in accordance with the above method as disclosed in the parent case) of live voice or text chat in real-time dynamic fashion. Spoken conversations may be automatically profiled using acoustic speech-to-text methods. In accordance with the methods for automatic construction of virtual communities (in this case for dynamic chat) as disclosed in the parent case, these communities can be constructed automatically and on-the-fly based upon similarities of chat content, queries, navigational (content) selection from automatically constructed menu and/or user profile similarity or attributes of the profiles of the users which may be selected. Thus, in the former case a user could in the case of a pull down menu gain access by title of key phrase (which are ideally automatically and dynamically created and updated in accordance with the content of dynamic spoken or recorded conversation spoken or typed in real time and which may be further dynamically categorized and re-categorized within a hierarchical menu structure format. Although, the



1 word(s)/phrase(s)(again, so long as the user's profile is within the disclosure policy of the  
2 transmitting user. If a willing recipient so desires for both the chat and instant messaging  
3 variations, s/he may receive a copy of the text or voice transmission or if the user is not able  
4 to access it in real-time or otherwise. Conversely, the user may as a sender wish to send the  
5 transmission to

6 1. All users who have indicated potential interest in the contents (and/or profile or identity  
7 of the sender) or users who match a particular user profile (e.g., interest or proficiency such  
8 as which is of measured similarity to the contents of the transmission) and/or

9  
10 3. Users who are in the process of engaging in a user-user (or multi-user) dialog, which is of  
11 measured similarity to that of the transmission.

12  
13 An example application of the present system includes, for TV viewers, the ability to gain  
14 access to different types of feedback from other users who are also presently viewing or had  
15 previously viewed a particular TV program, e. g., a comedy, a news story or political speech  
16 in which the present technique could be used effectively as a filter allowing certain types of  
17 user or user attributes to reach the user and others to be suppressed or squelched. In the  
18 latter example, a user who considers herself to be a female liberal may provide settings  
19 during a campaign speech by a conservative Republican speaking out against abortion to  
20 listen to acoustic/verbal gestures by other liberal females advocating abortion). Audible  
21 comments may be unfiltered and heard by the user, if desired.

22 If the program is a re-broadcast, more elaborate filtering capabilities are conceivable  
23 involving statistical analysis of the spoken language content (such could be conceivably  
24 performed for real time live information, however, at the expense of a slight delay). As  
25 suggested, it may be also possible to observe textually (e.g., through speech to text  
26 techniques) or hear spoken conversations as they exist between individuals possessing  
27 desired attributes, containing content characterized by key words or phrases from a menu  
28 selection or provided by/between a desired individual or individuals respectively.

29  
30 In another application, users viewing video or streaming video content may upon their  
31 permission, agree to disclose their user profiles and be acoustically monitored such that

1 comments, verbal, verbal/audible gestures and expressions and/or video of Effectuated thereby  
2 may be of interest to users

3  
4 In conjunction with an educational program a user may wish to observe comments by  
5 those other individuals who are extremely knowledgeable in the field (to collect  
6 supplemental information and other feedback about the presently viewed content).

7 Similarly, in the political speech example the user may desire to hear the gestures and/or  
8 spoken comments, reactions and/or brief dialogues occurring during the course of the  
9 political speech by individuals sharing the same political views as the user and which are  
10 considered to be part of the educated elite. In another variation, individuals who  
11 represent the exemplars of different cluster of users (and/or those which are manually  
12 selected as providing interesting and/or entertaining feedback) could be heard by the user  
13 collectively and the reactions expressed as acoustic feedback could be statistically  
14 analyzed in dynamic fashion and presented to the user as a breakdown of the user  
15 attributes which presently characterize the present reaction or response of the virtual  
16 audience.

17  
18 In a future application, such system could be extended to such things as immersive virtual  
19 reality (stationary or non-stationary) systems and/or video gaming systems in which the  
20 reactions of characters (or even different "personals") to various situations and events  
21 could be developed from data collected from the reactions of actual users comprising the  
22 different user clusters as above suggested.

23  
24 Applications to Constructing User Profiles and Matching Users by Similarities in Their  
25 Social and Psychological Profiles and/or Life Circumstances and Experiences

26  
27 There is an untapped opportunity with potentially considerable and deep implications  
28 which could be used to enable very detailed assessment and associated profiling of  
29 individuals. The parent patent describes a very comprehensive methodology by which  
30 users could potentially be profiled so as to define their preferences across almost any kind  
31 of content or commercial products and services. Surveys used to reveal psychological

09699098-102700

1 traits are also used. In the present extension, we first suggest a primary system  
2 methodology in which extremely detailed information is collected and aggregated into  
3 the user profile particularly pertaining to the user” psychological profile (user  
4 preference/interest information demographics, etc. are also useful data in that there are  
5 likely correlations which exist with certain psychological attributes of the user). In the  
6 clinical field of psychoanalysis, much information regarding the psychological and  
7 psycho-pathological characteristics of the user is determined, however in compliance  
8 with the confidentiality requirements of the doctor-patient privilege, most if not all of this  
9 information remains isolated and never accessible by a large-scale statistical database for  
10 purposes of cross-correlation of psychological characteristics, phenomena, psycho-  
11 pathology (as well as other clinical pathology), as well as behavior, interests, preferences  
12 and more specifically identifiable behavior patterns as well as notable life experiences.  
13 Part of the answer to this dilemma may well lie in the application of secure data  
14 disclosure to SDI, the pseudonym proxy server and (for extra security) randomized  
15 aggregates for purposes of harvesting the statistical value within the data and (if needed),  
16 interacting with the user while maintaining completely secure and private individual user  
17 data. The present system suggests that there may be much greater accuracy which could  
18 be achieved in the profiling of individuals for use in a number of applications such as the  
19 psychological modeling of patients for use in the diagnosis, treatment, prediction  
20 (including predictive treatment) of psycho-pathology, the matching of “similar” patients  
21 together (for group therapy and/or pseudonymous electronic or physical mail  
22 correspondences) for providing support, inter-patient communication, counseling, etc.  
23 (for which present traumatic life experiences, if relevant, may be quite applicable as  
24 well), as well as the extended/improved modeling/prediction of user interests/preferences  
25 enabled by SDI.  
26 Another correlated data type could be employee file databases through which correlations  
27 and predictions of employee behavior and performance, recommendation as to remedies  
28 for existing problems and performance enhancing tactics again could be achieved.  
29 Unfortunately not all data is likely to be accessible for all users (e.g.,  
30 browsing/transaction histories are likely to be available but often not information  
31 regarding a user’s life experiences or psychological history). To fill in this missing data,





1 example, disclose to the employer "Mr..X has a psychological profile feature which is  
2 highly correlated with a high aptitude for the present job which s/he is applying for and  
3 past job experience Y is 50% more likely than the average employee for that position to  
4 lead to a subsequent promotion and several "employee of the month awards"

5  
6 Application to Identifying Similar Individuals Behaving Under Similar Circumstance

7  
8 The present user profile information could also be extremely relevant for a novel Internet  
9 based service by which individuals could pseudonymously release their profiles and/or  
10 define specific situations and/or conditions (which could be actual, hypothetical or  
11 combination thereof) and the system will identify other scenarios (of like pseudonymous  
12 individuals and/or circumstances) to the general scenario and emphasizing and  
13 prioritizing certain particular aspects thereof. Examples include identifying individuals  
14 who had previously (or are presently) in the process of making a very similar professional  
15 decision, interpersonal decision.(such as regarding a present or a potential spouse).  
16 Ideally, such interpersonal decisions would involve a very similar situation (such as a  
17 decision), a very similar individual in the analogous position of the user (and ideally, if  
18 possible other similar party(s) fulfilling the analogous role(s) as the present situation  
19 which confronts the user. The relative importance (priority) of the degree of similarity of  
20 various analogous roles in an analogous situation is determined in part by the type of  
21 information which the user would like to determine. For example, determining what a  
22 particular individual would do in a similar circumstance may depend upon how  
23 dependent that user's action .or decision is upon the circumstances and conditions versus  
24 another individual(s) (and which individual(s) and/or the two (or more) individuals  
25 involved to which the context of the situation primarily relates. Of course, the nature of  
26 the action in question may predicate the relative importance of one or more of the above  
27 factors. These determinations are often quite subjective in nature. Requestees may also  
28 (if the choice exists) be situations which are either of a historical nature (i.e., the situation  
29 had occurred in the past) or of a present nature. In some cases in the latter case, is  
30 preferable as the user may actively request, e.g., the party to which s/he directly relates or  
31 is potentially affected by to provide advice and/or feedback or to what s/he would do

under a particular "what if" scenario which could be very specific and relevant to the present circumstances of the user. In a novel variation of the method, the party to which the user relates is commissioned by the user to do or say something in order to receive a reaction or response. Depending upon how similar the situation actually is, the prompt could be very similar or partially relevant to the action contemplated by the user in his/her own analogous situation. Or the user may wish to collect feedback from analogous individuals in similar situations regarding several hypothetical actions (perhaps acted out among several different groups) resulting from decisions in order to enhance the user's information about the user(s) and the associated situation. Alternatively, historical cases could include requests for valuable advice (as the experiences and the associated consequences are more likely to have been lived out). Typically it is most valuable (if a number of similar circumstances exist ) to collect multiple reactions/responses with which to form a statistical basis for the information. The biggest problem with the present system is collection of relevant and sufficiently detailed data relating to the relevant circumstances which must be identified and accessed. Psychotherapeutic (clinical) databases would be one source, as would be spoken or typed synchronous or asynchronous communications between individual as well as (the emerging) ubiquitous computing environments in which users' actions in spoken communications could be potentially monitored persistently off-line as well as on-line (for which SDI's privacy-enhanced profiling architecture would be particularly appropriate). Certainly, if the circumstances affecting the user (requestor) are of a long-term nature, periodic informational updates could be provided to him/her and/or the mutual sharing and advising of the similar experience could be the basis for a two-way exchange of information between the parties on a short or long-term basis.

#### AUTOMATED LEARNING OF USER BEHAVIOR AS PART OF AN ENHANCED MEETING SCHEDULER AND CORRESPONDENCE FILTER

The present section provides application level extensions to similar inference-based rules recommendation and generation functionality such as is described in conjunction with the Smart Home description , smart office as well as the learning based methods disclosed

1 within the section entitled "Resolution credentials" in the access reachability introduction  
2 and profile disclosure between two or more parties as disclosed

3  
4 In order to utilize LEIA as a user agent, which is integral and ubiquitous within the work  
5 environment, it is necessary to make its learning as seamless and invisible from the user's  
6 perspective as possible.

7  
8 The next level in applying the capabilities of LEIA is in being able to implicitly learn the  
9 context of the user's present temporal activities as well as establish relative priorities of  
10 the activities that the user (and others who intend to interface with the user) are engaged  
11 in and from the user's previous behavior to similar conditions automatically generate  
12 appropriateness functions (or rules) to automate the process of handling meeting  
13 schedules and filtering real time requests for correspondence with the user. It is possible  
14 to infer:

- 15  
16 1. The content profiles which are associated with the user's present activities.  
17  
18 2. (Often obviously) The nature of the user's present activities by observing what sorts of  
19 actions the user is engaged in, in the office.

20  
21 These content profiles and activity indicators may be passively observed by the user's  
22 present interactions (and timing thereof) with his/her PC, other smart appliances and data  
23 and voice communications which the user engages in (telephonically or from standard  
24 face to face dialogues) e.g., suggesting when the user is reading the morning news,  
25 checking phone messages, responding to e-mail, engaged in a particular project, etc. If  
26 these clues are not observable (or not positively identifiable) cyclical time dependent  
27 patterns may be used as implicit indicators of the user's present activities (see the above  
28 section entitled "Further Applications of LEIA"). Alternatively, the system could prompt  
29 a request of confirmation of the user's present activities and possibly on clues such as  
30 verbal regarding present and future established or changing activities and plans These  
31 same methods for utilizing time pattern elements which are used to predict the location of

1 the user may be readily extended to predict or help to substantiate other statistical user  
2 variables relating to the types of actions and content related profiles characterizing the  
3 user's temporal behavior patterns. The following application of these techniques is used  
4 to enhance the automatic meeting scheduler:

5  
6 The identity of the task or request, the activity and content attributes of the task or request  
7 profiles (target object profiles) may be developed utilizing the priority of that task  
8 relative to the user. This priority is estimated from the priority of the task/request (and its  
9 associated attributes) over other tasks/requests by the user. In order to better substantiate  
10 the relative importance to the user as inferred regarding these activities, the system  
11 should also display the scalar ratings which the user could adjust manually.

12 Nonetheless, there may be a degree of uncertainty in the user's intentions and LEIA's  
13 inference of the user's intentions. Moreover, unanticipated changeability by the user  
14 (uncertainty) may occur perhaps with increased frequency under certain variables like the  
15 identity of the requestor, the time of day, context of the users other present activities or  
16 when other meetings and obligations are pending which possess a relatively greater  
17 priority with the user (i.e., if they often occur spontaneously). In order to take this  
18 uncertainty factor into consideration, an overall statistical estimate may then be  
19 calculated taking into account the predicted statistical probability of each party ultimately  
20 being available for the meeting. This serves the purpose of both refining the automatic  
21 scheduling process and informing would be visitors before hand as to the relative  
22 importance that their prospective meeting is (or increasingly becomes as the meeting time  
23 approaches) to the employee as well as the overall estimated degree of certainty of the  
24 meeting. This feature may be particularly useful for impromptu visitors where  
25 considerable uncertainty exists as to whether and to what degree the visitor is imposing  
26 upon the employee's schedule.

27  
28 The temporal context of user behavior (using temporal time series analysis in conjunction  
29 with use of content analysis and user behavior assessment may be both important  
30 predictors as to whether a prospective meeting of a disclosed purpose is worthwhile for  
31 the employee based on present/future priorities and at which most likely times. By far,

1 its most valuable application , however is in its ability to perform content analysis and  
2 user behavior assessment in real time in order to enhance LEIA's ability to determine  
3 whether and to what degree a would-be unannounced visitor or telephone or intercom  
4 callers unannounced contact is relevant to the employee" present activities and if not to  
5 coordinate with the visitor's agent the time or times when it would be most relevant to the  
6 employees anticipated activities for example there are certain activities during which an  
7 employee doesn't wish to be usually disturbed, e.g. while reading the morning news,  
8 taking a coffee break, calling home, responding to an important email message, ten  
9 minutes before an important meeting. Often within an organization individuals will  
10 attempt to speak with one another as such , the intended duration affects priority as does  
11 the identity of the individual in the organization. Often this variable can be automatically  
12 predicted by LEIA . In such a situation, the relative priority of the user's preexisting  
13 priority tasks is weighted against the priority of the requested tasks prior to notification.  
14 The above application can be readily extended as well to phone calls, in which the user  
15 could be automatically prompted by LEIA to disclose his/her purpose of the call and/or  
16 his/her identity if necessary. This application would be a useful extension to the  
17 telephony variation of the email filter described in issued patent entitled "System and  
18 Method for Electronic Identification of Desirable Objects in which the above additional  
19 inputs about employees and visitors may provide additional useful feedback in  
20 automatically determining the rules dictating how to best handle the call or forward it to a  
21 more appropriate employee on behalf of the employee. If there is a certain degree of  
22 uncertainty in the agent's priority determination or assessment of the task affected the  
23 most appropriate rule the employee may observe the rule just prior to automatic  
24 implementation. Or the action which it represents. The present value of the contextual  
25 employee activity assessment method as disclosed, if extended to advertising could be  
26 significant. For example the knowledge of the context of the user's present activities  
27 could determine what type and when an advertisement targeted to an employee, e.g. a  
28 decision maker will most likely elicit the best response (where the advertisement  
29 matches the professional interest profile) is an ad which corresponds with particular  
30 activities of the user ( e.g., which relate in some way with the ad/promotion. As  
31 discussed direct voice telephony and email dialogues may provide very useful clues

